Universitatea
Transilvania
din Brașov

HABILITATION THESIS

Title: Applied Intelligence - Machine Learning Across Interdisciplinary Technologies and Systems

Domain: Computer Science

Author: Assoc. Prof. Dr. Alexandra Băicoianu

University: Transilvania University of Brașov, România

BRAŞOV, 2025

Contents

Acknowledgement

I would like to express my sincere gratitude to all the collaborators from *Transilvania University* of Brașov - both from the Faculty of Mathematics and Computer Science and from the research teams of AI4AGRI, AI4RISK, SEEN and IMINT. Their continuous support, expertise, and shared enthusiasm have been significant in shaping many of the research outcomes included in this thesis. Most of them are also co-authors of the publications referenced throughout this work, and their contributions have been truly invaluable.

I am also deeply grateful to my colleagues at *Siemens Industry Software*, whose ongoing support and encouragement have played a key role in my academic journey. Through our joint involvement in impactful and challenging research projects, such as DITARTIS, NEMOShip, IMOCO4.E, INVENT, DILSIMEV, PANDA, and URBIVEL, they have not only offered me professional growth but have also actively supported the university-industry partnership over the years. I am fortunate to now collaborate closely and enthusiastically with several former students (now my colleagues) whose professional growth and dedication are a continuous source of inspiration. The bachelor's and master's theses developed in partnership with *Siemens Industry Software* have also played a key role in consolidating many of the ideas and directions presented in this thesis.

I further thank all my colleagues from *Transilvania University* who have supported me in both teaching and research activities. I am especially grateful to those who went above and beyond by helping me stay focused and consistent throughout this thesis - revisiting experiments, reviewing drafts, and providing thoughtful corrections and feedback. Their dedication and collegial spirit made a real difference in the quality of my work. Last but certainly not least, I am profoundly thankful to my family and all those dear to me, whose encouragement and presence continue to provide the energy and motivation I need every day.

As Kurt Lewin once said, *There is nothing so practical as a good theory.* This idea has guided much of the work presented in this thesis, where theoretical insights have often proven to be powerful tools in addressing real-world challenges.

A Rezumat

Această teză de abilitare, centrată pe tema fundamentală *Inteligență Aplicată - Învățare Automată în Tehnologii și Sisteme Interdisciplinare*, examinează aprofundat modul în care progresul rapid în inteligența artificială, învățarea automată și analiza datelor redefinește continuu abordările esențiale în inginerie, științele mediului și diverse domenii industriale aplicate. Această lucrare reunește o serie de contribuții de cercetare care explorează integrarea tehnicilor de AI în sisteme reale, cu accent pe aplicații cu impact și relevanță interdisciplinară ridicate.

Lucrarea este structurată în trei capitole conectate tematic:

1. **Inteligența artificială în observarea Pământului și monitorizarea agricolă** - acest capitol abordează utilizarea modelelor bazate pe analiză de date pentru clasificarea tipurilor de utilizare a terenurilor, monitorizarea culturilor și evaluarea mediului pe baza imaginilor satelitare. Se concentrează pe provocări precum vizualizarea inteligentă și interpretarea indicilor de vegetație, integrarea datelor provenite din surse multiple și crearea de seturi de date pentru o clasificare precisă a culturilor, precum și pe extragerea de informații utile prin studii de corelație în monitorizarea agricolă bazată pe observarea Pământului.

   Numeroase dintre contribuțiile prezentate în acest capitol au fost dezvoltate și rafinate în cadrul unor proiecte de cercetare semnificative. De exemplu, mai multe studii au servit drept baze esențiale pentru abordarea obiectivelor proiectului european *AI4AGRI Romanian Excellence Center on Artificial Intelligence for Earth Observation Data in Agriculture*. În calitate de cercetător afiliat *Universității Transilvania*, rolul autorului a implicat procesarea imaginilor multi- și hiperspectrale utilizând metodologii de învățare automată pentru agricultura de precizie, activitate desfășurată frecvent în colaborare strânsă cu echipa proiectului AI4AGRI. Echipa a fost distinsă cu un premiu în cadrul programului "Horizon Europe Research Teams Awards" organizat de UEFISCDI, evidențiind contribuția sa semnificativă la creșterea vizibilității României în programele europene de cercetare.

2. **AI pentru monitorizare, diagnostic și optimizare în sisteme complexe** - acest capitol se concentrează pe dezvoltarea de soluții bazate pe inteligență artificială pentru supravegherea condiției, detectarea defectelor și optimizarea sistemelor, cu aplicabilitate în domenii industriale, software, medicale și maritime, având ca scop principal creșterea performanței și întreținerea predictivă. Această abordare inte-

grează o serie de direcții specifice, precum monitorizarea condiției și gestionarea defectelor asistate de inteligență artificială în sisteme industriale și software, aplicarea algoritmilor evolutivi pentru optimizarea sistemelor structurale, utilizarea inteligenței artificiale în triajul medical și suportul decizional clinic, optimizarea acurateței în prognoza seriilor temporale, dar și inovația digitală și optimizarea proceselor în domeniul maritim.

O parte semnificativă a contribuțiilor prezentate în acest capitol provin din proiecte de cercetare interdisciplinare. Aceste proiecte, incluzând DITARTIS, NEMOShip și IMOCO4.E, au necesitat în mod specific integrarea componentelor de inteligență artificială pentru diverse obiective de optimizare și automatizare. Această abordare multidisciplinară nu numai că a facilitat dezvoltarea de metodologii inovatoare, dar a demonstrat și aplicabilitatea practică și impactul soluțiilor propuse în diverse domenii.

3. **Modele computaționale și comportamente inteligente** - acest capitol investighează automatele celulare și, in particular, automatele celulare neuronale ca modele computaționale puternice care demonstrează comportamente inteligente într-o varietate de aplicații. Acoperă capacitatea lor de a învăța să genereze structuri complexe și texturi diverse, precum și eficacitatea lor în modelarea predictivă pentru sisteme dinamice, cum ar fi răspândirea bolilor. În mod specific, acest capitol explorează înțelegerea dinamicii de învățare și creștere în cadrul automatelor celulare neuronale și dezvoltarea automatelor sensibile la semnal, utilizate pentru sinteza versatilă de multitexturi.

   Deși cercetarea prezentată în acest capitol nu este direct legată de rezultate specifice obținute în cadrul unor proiecte externe, ea reprezintă o explorare aprofundată și continuă a aspectelor fundamentale ale inteligenței artificiale neuronale și ale paradigmelor computaționale avansate. Aceasta investighează modul în care modelele computaționale pot manifesta comportamente inteligente și se pot adapta eficient la sisteme dinamice.

Cercetarea prezentată în aceste capitole este fundamentată pe peste 40 de studii distincte, la care autorul a contribuit fie ca autor principal, fie ca și coautor. Fiecare capitol se bazează pe descoperiri care sunt validate suplimentar prin cel puțin o publicație într-un jurnal clasificat Q1 și susținute de articole adiționale din jurnale Q2/Q3, precum și de numeroase lucrări publicate în cadrul conferințelor de specialitate. Toate lucrările citate sunt direct aliniate cu subiectele specifice și întrebările de cercetare abordate în teză, asigurând atât rigoare academică, cât și coerență tematică.

În ansamblu, aceste contribuții subliniază valoarea practică și profunzimea teoretică a metodelor de inteligență artificială în abordarea provocărilor diverse și în continuă evoluție, de la detectarea defecțiunilor mecanice și modelarea bolilor la agricultura inteligentă și mobilitatea sustenabilă.

Această teză își propune nu doar să consolideze rezultatele cercetării aplicate, ci și să formuleze o viziune coerentă, că puterea transformatoare a inteligenței artificiale constă în capacitatea sa de a permite luarea deciziilor inteligente, reziliente și conștiente de context în orice domeniu științific și/sau ingineresc. Această explorare aprofundată surprinde esența *Inteligenței Aplicate*, evidențiind cât de larg răspândită a devenit învățarea automată ca un instrument esențial în diverse tehnologii și sisteme interdisciplinare. Prin conectarea perspectivelor bazate pe date cu aplicațiile practice, aceasta facilitează progrese în domenii precum ingineria, medicina, observația terestră și grafica computerizată, unde automatizarea inteligentă și recunoașterea tiparelor sunt esențiale pentru rezolvarea unor provocări complexe.

Summary

This habilitation thesis, which is built upon the central theme of *Applied Intelligence - Machine Learning Across Interdisciplinary Technologies and Systems*, thoroughly examines how the rapid progress in artificial intelligence (AI), machine learning (ML), and data analytics is continuously reshaping fundamental approaches in engineering, environmental sciences, and various applied industrial fields. This work brings together a series of research contributions that explore the integration of AI techniques in real-world systems, with a focus on high-impact applications and cross-domain relevance.

The work is structured into three thematically connected chapters:

1. **AI in Earth Observation and agricultural analytics –** addressing the use of data-driven models for land use classification, crop monitoring, and environmental analysis based on satellite imagery, including challenges related to intelligent visualization and interpretation of vegetation indices, the aggregation of multisource data and development of datasets for accurate crop classification, and the extraction of actionable insights through correlation studies in Earth Observation-based agricultural monitoring.

   Many of the contributions within this chapter were developed and refined within the framework of significant research projects. For instance, several studies served as essential foundations for addressing the objectives of the AI4AGRI Romanian Excellence Center on Artificial Intelligence for Earth Observation Data in Agriculture. As a researcher affiliated with *Transilvania University*, the author's role involved processing multi- and hyperspectral images using machine learning methodologies for precision agriculture, often in close collaboration with the AI4AGRI project team. This team was notably recognized with an award under the 'Horizon Europe Research Teams Awards' program by UEFISCDI, highlighting its impact on increasing Romania's visibility in European research programs.

2. **AI for monitoring, diagnosis, and optimization in complex systems** - developing AI-driven solutions for condition monitoring, fault detection, and system optimization across industrial, software, medical, and maritime applications, focusing on performance enhancement and predictive maintenance. This contains specific problematics such as AI-driven condition monitoring and fault management in industrial and software systems, the application of evolutionary algorithms for structural system optimization, AI for medical triage and clinical decision support, the optimization of

accuracy in environmental time series forecasting, and digital innovation and optimization in maritime applications.

Many of the studies presented in this chapter originate from interdisciplinary research projects. These projects, including DITARTIS, NEMOShip, and IMOCO4.E, specifically required the integration of AI components for diverse optimization and automation goals. This multidisciplinary approach not only facilitated the development of innovative methodologies but also demonstrated the practical applicability and impact of the proposed solutions across various domains.

3. **Computational models and intelligent behaviors** - investigates Neural Cellular Automata (NCA) and Cellular Automata (CA) as powerful computational models that demonstrate intelligent behaviors across various applications. It covers their capacity for learning to generate complex structures and diverse textures, as well as their effectiveness in predictive modeling for dynamic systems, such as disease spread. Specifically, this chapter explores understanding the learning and growth dynamics within NCAs and the development of signal-responsive NCAs for versatile multi-texture synthesis.

   While not directly tied to specific external project acknowledgements, the research in this chapter represents a deep and continuous exploration into the foundational aspects of neural AI and advanced computational paradigms, investigating how computational models can manifest intelligent behaviors and adapt to dynamic systems.

The research presented across these chapters is grounded in over 40 distinct studies, to which the author has contributed either as lead or co-author. Each chapter builds upon findings that are further validated through at least one publication in a Q1-ranked journal, and supported by additional articles from Q2/Q3 journals, as well as numerous peer-reviewed conference papers. All referenced works are directly aligned with the specific topics and research questions addressed in the thesis, ensuring both academic rigor and thematic coherence.

Collectively, these contributions highlight the practical value and theoretical depth of AI methods in addressing diverse and evolving challenges, from mechanical fault detection and disease modeling to smart agriculture and sustainable mobility.

This thesis aims not only to consolidate applied research findings, but also to articulate a unifying vision: that the transformative power of AI lies in its ability to enable intelligent, resilient, and context-aware decision-making across any scientific and/or engineering domains. This in-depth exploration captures the essence of *Applied Intelligence*, highlighting how widely machine learning has become an essential tool across various technologies and interdisciplinary systems. By bridging data-driven insights with practical applications, it enables advancements in fields such as engineering, medicine, Earth observation, and computer graphics, where intelligent automation and pattern recognition are key to solving complex challenges.

B Scientific and professional achievements and the evolution and development plans for career development

(B-i) Scientific and professional achievements

Introduction

I hold a Bachelor's degree in Mathematics and Computer Science from the Faculty of Mathematics and Computer Science, *Transilvania University* of Brașov, graduating with a GPA of 9.53. In 2008, I completed my Master's degree in "Algorithms and Software Products" at the same institution with a perfect GPA of 10.00. Beyond my academic qualifications, I have also pursued professional development in education and training. I completed the Teacher Training Department coursework at *Transilvania University of Braşov* and the trainer certification program within the project "Development of Specific Competencies for Designing and Implementing Distance Education Technology Training Programs". Following these achievements, I pursued doctoral studies and earned my Ph.D. in Computer Science from *Babeș-Bolyai University* of Cluj-Napoca in January 2017 (3148/30.01.2017). My doctoral thesis, titled *Some Other Issues in Discrete Optimization*, was defended on September 30, 2016, under the scientific supervision of Professor Militon Frențiu.

Over the course of 18 years in higher education, I have progressed through successive academic roles: collaborator (2007-2008), teaching assistant (2008-2015), assistant professor (2015-2017), lecturer (2018-2024), and currently hold the position of associate professor (2024-present). All these positions have been held at the Faculty of Mathematics and Computer Science, *Transilvania University* of Brașov.

As a professor, I have developed comprehensive theoretical and practical course materials for a diverse range of subjects including: Formal Languages, Compilation Techniques, Fundamental Algorithms, Object-Oriented Programming, Standard to Modern C++, and Digital Image Processing. The quality of my teaching has been consistently recognized through peer and student evaluations, predominantly receiving "excellent" ratings. A particular point of pride is that since the beginning of "the most appreciated professor" program at *Transilvania University*, I have consistently ranked among the top-rated faculty members in the Faculty of Mathematics and Computer Science.

While maintaining my academic commitments, I have also extended my expertise into the corporate sector. This parallel career path has included roles as C++ Trainer

at *Route66 SRL*, Brașov (2013-2015), and Software Engineer at *Siemens Industry Software*, Brașov (2017-present). Soon after joining Siemens Industry Software, I took on the additional responsibility of contributing to the research department, actively engaging in the preparation and submission of project proposals, including national and European funding programs. This involvement has allowed me to apply my academic background to real-world innovation challenges while collaborating in interdisciplinary teams.

My involvement with *Siemens Industry Software* and *Route 66* has provided invaluable insights into industry expectations of our graduates, enabling me to better guide and support students' development at the Faculty of Mathematics and Computer Science. The technical knowledge gained through these corporate experiences has directly influenced the continuous updating of laboratory and seminar curricula. Furthermore, this industry collaboration has led to the development of several specialized elective courses, including: *Standard to Modern C++*, *Modern C++ - Managing Networking Projects*, *C++ Modern Applications in Artificial Intelligence*, *Cloud Computing*, *Full Stack Development of Web-based ML Applications*.

My professional approach emphasizes effective teamwork and adaptability in challenging situations. Success in my role comes from maintaining professional integrity, strong analytical skills, and continuous self-development. I regularly update my teaching methods to incorporate modern technological tools and interactive systems. As part of my academic responsibilities, I mentor students and university assistants in their research endeavors, supervising various academic works including scientific presentations, bachelor's theses, and master's dissertations.

Throughout my career, I have maintained excellent communication and collaborative relationships with colleagues, attributed to my dynamic spirit, demonstrated perseverance, and commitment to objective evaluation. This collaborative approach, combined with my industry experience and academic expertise, has created a comprehensive educational environment that bridges theoretical knowledge with practical application, preparing students for both academic excellence and industry demands.

Throughout this multifaceted career, I have consistently focused on enhancing both my professional capabilities and students' learning outcomes, continuously improving pedagogical methods, and ensuring students acquire specific competencies in both theoretical and laboratory courses. My dedication to academic excellence is further demonstrated through numerous publications in national and international scientific conferences, as well as specialized books and courses, all aligned with the Ministry of Education's requirements and ARACIS standard evaluation systems.

Research projects

Throughout my career, I have been actively involved in numerous research projects, collaborating with/as part of teams at both *Transilvania University* of Brașov and *Siemens Industry Software*. These projects reflect my expertise across multiple domains, with par-

ticular emphasis on:

- Application of AI methods in interdisciplinary contexts
- Optimization techniques and algorithms
- Cloud platforms and technologies
- Data analysis and predictive modeling

For each project detailed below, I highlight my specific role (whether as project responsible person, *Siemens Industry Software* team member, or *Transilvania University* researcher), my direct contributions, and the project's broader impact. The projects demonstrate a consistent focus on practical applications of advanced technologies, particularly in areas where AI, optimization methods, and cloud computing intersect with industry needs and academic research.

Below are some notable ongoing and completed projects, organized to reflect both their current status and my specific involvement in each of them.

Category I - Ongoing projects

- **DITARTIS** *Digital Technologies and Artificial Intelligence (AI) solutions*
  - Details: https://ditartis.utcluj.ro/#project, https://cordis.europa.eu/project/id/101079242
  - Role: Responsible project manager within the *Siemens Industry Software* institution
  - Research focus: Participation in two integrated work packages: *RS3*: AI-based electrical machines and drives (design, analysis, control, testing) & *RS4*: Data-driven condition monitoring and predictive maintenance in EMS.

  Active participation in the mentorship program for young researchers, combined with engagement in research projects centered on data-driven condition monitoring and predictive maintenance within Energy Management Systems (EMS).

- **NEMOHIP** *New modular Electrical architecture and digital platforM to Optimise large battery systems on SHIPs*
  - Details: https://cordis.europa.eu/project/id/101096324
  - Role: *Siemens Industry Software* researcher
  - Research focus:
    - Development of predictive models for ship data analysis
    - Data analysis and predictive modeling

As part of this project, I am a member of the team that analyzed and processed historical data from ships, developing predictive models to provide informed recommendations regarding future changes, thereby ensuring continuous optimization of processes.

- **IMOCO4E** *Intelligent Motion Control under Industry 4.E*
  - Details: https://www.imoco4e.eu/
  - Role: *Siemens Industry Software* researcher
  - Research focus: Software solution development for synthetic data generation using Simcenter Amesim

  As part of this project, I contributed to the development of a software solution for synthetic data generation using Simcenter Amesim, thereby enhancing simulation and analysis processes.

- **AI4AGRI** *Romanian Excellence Center on Artificial Intelligence on Earth Observation Data for Agriculture*
  - Details: https://ai4agri.unitbv.ro/
  - Role: *Transilvania University* member
  - Research focus: AI applications in Earth Observation Data for Agriculture

  As part of this project, I am involved in the processing of multi- and hyperspectral images using machine learning methodologies, providing diverse solutions in the field of precision agriculture.

  The AI4AGRI project team was awarded under the "Horizon Europe Research Teams Awards" program by UEFISCDI, https://uefiscdi.gov.ro/index.php, Secțiunea II, *Listă Premiere Orizont Europa - Echipe de cercetare - dupa suplimentare buget*. This national initiative recognizes and rewards Romanian research teams that have successfully participated in Horizon Europe projects, aiming to increase Romania's visibility and participation in European research programs.

- **SOL-2024 - IMINT** *IMage INTelligence for the Black Sea*
  - Details: PN-IV-P6-6.3-SOL-2024-0124, Contract no. 21Sol(T21)/2024, https://uefiscdi.gov.ro/sol-2023-imint-pentru-marea-neagra-frontiere-mine
  - Role: *Transilvania University* member
  - Research focus: AI-powered semantic interpretation of remote-sensing and EO data for Black Sea threat detection

  A national Romanian project coordinated by the National University of Science and Technology Politehnica Bucharest, in partnership with Military Technical Academy

Ferdinand I, Romanian Space Agency, Naval Academy Mircea cel Bătrân, Intergraph Computer Services SRL, and *Transilvania University* of Brașov. The project implements cutting-edge AI tools for advanced image interpretation and security threat detection in the Black Sea region.

- **AI4RISK** *Platform for Fusion and Management of Multi-source Data Collections*
  - Details: https://ai4risk.ro/, 3Sol(T3)/02.09.2024, PN-IV-P6-6.3-SOL-2024-2-0251
  - Role: *Transilvania University* member
  - Research focus: Multi-source data fusion and AI-based risk analysis

  This national Romanian project focuses on developing advanced AI tools for multi-level analysis of diverse data sources. AI4RISK proposes an innovative platform to strengthen national security by integrating multi-source data, processing real-time image streams, and leveraging explainable artificial intelligence (XAI) for risk analysis and assessment.

- **SEEN** *Social Entrepreneurship Ecosystem Network*
  - Details: https://www.unitbv.ro/cercetare/rezultatele-cercetarii/, Erasmus+ KA220-HED cooperation partnership in higher education
  - Role: *Transilvania University* member
  - Research focus: Entrepreneurial education and AI-supported business planning

  This project aims to establish a comprehensive network for training future research and education professionals in entrepreneurship. It combines practical workshops and teaching activities with innovative AI tools to support business plan evaluation and enhance students' innovation and creativity capabilities. This project involves international partnership including institutions from Bulgaria, Lithuania, Netherlands, Romania, and Spain.

- **Post-quantum Digital Signature using Verkle Trees**
  - Details: https://quantum.unitbv.ro, NATO SPS G7394, 2024-2026
  - Role: *Transilvania University* member
  - Research focus: Designing and implementing a post-quantum digital signature scheme that combines Verkle trees and advanced random number generation for secure, efficient, and robust cryptography.

  The project develops a novel post-quantum digital signature scheme using Verkle trees, polynomial commitments, and vector commitments to enhance security and

efficiency. It integrates pseudo-random and quantum random number generators for greater robustness, addressing emerging quantum threats to cryptography and ensuring secure digital communications.

- **INVENT** *Innovative Multi-Life Battery Management System for future energy storage solutions*
  - Details: PN-IV-P7-7.1-PED-2024-1479, National research project in Climate, Energy and Mobility domain
  - Role: *Siemens Industry Software* researcher
  - Research focus:
    - Advanced battery management systems for second-life EV batteries
    - Cloud development and ML algorithms team member

  This project develops an innovative, next-generation Battery Management System (BMS) for repurposed Electric Vehicle batteries, leveraging a distributed architecture with a novel power line carrier protocol (xDCLink) for resilient, quasi-wireless monitoring. As part of the cloud development team, my responsibilities include implementing critical functionalities such as user authentication, data storage and streaming, security protocols, Battery Management Unit software update mechanisms, and machine learning algorithms for processing battery behavior data. This agnostic and versatile system addresses key energy storage challenges, aligning with EU Battery Regulation and Green initiatives.

Category II - Completed projects

- **HEART** *High-PErformance Computing of PersonAlized CaRdio-Vascular ComponenT Models*
  - Details: https://heart.unitbv.ro/, PN-II-PT-PCCA-2011-3.2-1229
  - Role: *Transilvania University* member
  - Research focus: Computational modeling for coronary artery disease

  This project focused on improving the diagnosis and management of coronary artery disease (CAD) through the use of advanced computational models. The main objective is to develop accurate and personalized methods for diagnosing and predicting the disease at the individual level, using detailed information about the anatomy and functionality of the coronary arteries.

- **DILSIMEV** *Driver-In-the-Loop SIMulation for safety-critical testing scenarios of Electric Vehicles*
  - Role: *Siemens Industry Software* researcher

  - Details: https://www.dilsimev.ro/, PN-III-P2-2.1-PED-2019-4366
  - Research focus: Autonomous driving application development

  The project aims to develop an advanced simulation tool to enhance the driving experience and increase vehicle safety in the new era of semi- and fully autonomous cars. Its outcome will be a novel human-machine interaction paradigm, implemented through a real-time decision-making system that integrates drivers' actions and personal need - including acquired and automated competencies - into the learning process of autonomous systems. Within this project, I was actively involved in WP5, which focused on analyzing, segmenting, and labeling driver actions in the Driving Simulator, contributing to the improvement of the autonomous driving application

- **PANDA** *Powerful Advanced N-level Digital Architecture for models of electrified vehicles and their components*
  - Details: https://project-panda.eu/
  - Role: *Siemens Industry Software* researcher
  - Research focus:
    * Development of Panda Explorer (cloud-based model/data storage)
    * Creation of Panda Model Converter
    * Implementation of cloud communication between Amesim and Typhoon

  Within this project, I was involved in the development of several key applications: Panda Explorer (a tool for cloud-based model/data storage), Panda Model Converter (a tool for transferring models from Amesim to Typhoon), as well as the component responsible for enabling cloud communication between Amesim and Typhoon.

- **URBIVEL** *Dezvoltarea unei platforme HIL pentru testarea bateriilor unui vehicul urban electric*
  - Details: http://urbivel.utcluj.ro/, Contract: 2/20.07.2017, https://www.epe.tuiasi.ro/2020/index_files/URBIVEL%20at%20EPE%202020.pdf, Proiect cofinanţat din Fondul European de Dezvoltare Regională prin Programul Operaţional Competitivitate 2014–2020 "Investiţii pentru viitorul dumneavoastră"
  - Role: *Siemens Industry Software* researcher
  - Research focus: Development of a Hardware-in-the-Loop testing platform for electric vehicle batteries

  I contributed to the development of a Hardware-in-the-Loop testing platform for the integration of energy sources in vehicles.

Professional activities and scientific leadership

I actively participate in various organizational committees for national and international scientific events, demonstrating my commitment to academic leadership and scientific community development.

Category I - Conference organization and scientific committees

- Scientific committee member, 36th International Business Information Management Conference (IBIMA), 2020
  https://ibima.org/conference/36th-ibima-conference/#ffs-tabbed-13
- Organizing committee member, The International Congress of Automotive and Transport Engineering (CONAT 2024)
  https://conat.ro/index.php/conat/2024/about/organizingTeam
- Committee member, MACOS, Brașov
  https://mateinfo.unitbv.ro/ro/macos.html/
- Workshop organizer, IEEE International Conference on Emerging Technologies and Factory Automation (ETFA 2023), 28th Annual Conference of the IEEE Industrial Electronics Society
  https://2023.ieee-etfa.org/main/static/files/ws_cfps/WS02_IMOCO4E.pdf
- Technical program chair and Track chair of "AI and Advanced Analytics in Energy Systems" session, FES 2025 Conference
  https://fes2025.com/conference-committees/#toggle-id-2
- Member of the organizing team (Scholarship chair), ACM womENcourage 2025
  https://womencourage.acm.org/2025/index.php/our-team/

Category II - Editorial activities and academic mentoring

I serve on the Editorial Board of the Editorial board member of the Bulletin of *Transilvania University* of Brașov, Mathematics and Computer Science section,
https://webbut.unitbv.ro/index.php/Series_III/Editorial_Board.
I have actively mentored three PhD candidates at *Transilvania University* of Brașov providing guidance on their research and thesis development: Iulian Popa (PhD in Computer Science, title of the thesis "Metaheuristic algorithms for optimization and their applications"), Bogdan Petre (PhD in Computer Science, title of the thesis "GPU-based methods in Bioinformatics"), and Ioana Lăzărescu (PhD in Mathematics, title of the thesis "Applications of differential geometry in machine learning").

Scientific publications

My research contributions comprise more than 40 articles indexed in Web of Science (WoS)/ISI and other international databases. After obtaining my PhD, I have published research papers as co-author in various scientific journals and conference proceedings, categorized according to their ranking as follows:

The Q1-ranked journals are:

- *Scientific Reports*, Springer Nature (Nature Portfolio), IF 3.9
- *Engineering Applications of Artificial Intelligence*, Elsevier Science, IF 8
- *IEEE Transactions on Geoscience and Remote Sensing*, Institute of Electrical and Electronics Engineers, IF 8.6, AIS 1.64
- *Big Earth Data*, Taylor & Francis Online, IF 3.8
- *Advanced Quantum Techologies*, Wiley Online Library, IF 4.3
- *Mathematics*, Multidisciplinary Digital Publishing Institute, IF 2.2
- *Big Data and Cognitive Computing*, Multidisciplinary Digital Publishing Institute, IF 4.4

The Q2/B-ranked journals and conferences are (selection):

- *IEEE Access*, Institute of Electrical and Electronics Engineers, IF 3.6
- *Heliyon*, Elsevier Science, IF 3.6
- *European Journal of Remote Sensing*, Taylor & Francis Online, IF 3.7
- *In Proceedings of the 18th International Conference on Software Technologies*
- *The 14 International Conference on Knowledge Science Engineering and Management*

The Q3/C-ranked journals and conferences are (selection):

- *The 17th International Conference on the Quality of Information and Communications Technology*
- *The 13th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*
- *The 28th International Conference on Emerging Technologies and Factory Automation*
- *Bulletin of the Transilvania University of Braşov Series III Mathematics and Computer Science*

- *International Conference on INnovations in Intelligent SysTems and Applications*
- *Big Data and Cognitive Computing*, Multidisciplinary Digital Publishing Institute, IF 4.4
- *Annals of the Tiberiu Popoviciu Seminar of Functional Equations*
- *The 21st International Conference on Applied Computing*
- *International Symposium on Signals, Circuits and Systems*
- *CSUM: Conference on Sustainable Urban Mobility*

The main directions of my research are clearly reflected both in the profile of journals and conferences where I have published, as well as in my current doctoral supervision activities. My research interests converge around several interconnected areas in intelligent systems and applications, each supported by different publications and practical implementations:

- AI Applications in interdisciplinary contexts
  Personal contributions in the field: [157], [30], [138], [15], [12], [70], [141], [86], [19], [106], [20], [131], [175], [134], [101], [133], [130], [132], [17], [21], [172], [33], [31], [11], [19], [174].
- Optimization techniques and algorithms
  Personal contributions in the field: [14], [30], [38], [16], [10], [18], [21], [17], [31], [175], [174], [122], [147], [13], [139], [32].
- Predictive modeling and data analysis
  Personal contributions in the field: [21], [141], [172], [19], [15], [12], [17], [86], [139].
- Intelligent condition monitoring using AI
  Personal contributions in the field: [15], [12], [70], [141], [86], [19].
- AI in geoscience and remote sensing applications
  Personal contributions in the field: [106], [20], [131], [134], [101], [133], [130], [132], [17].

These research directions are not only evident in my publication record but are also reflected in my academic involvement as a member of guidance committees for doctoral candidates at our doctoral school. This role allows me to contribute to the development of emerging researchers while ensuring the continuation and expansion of these critical areas of study. My integrated approach to research combines theoretical foundations with practical applications, demonstrating commitment to advancing innovations in intelligent systems across various domains.

Published books and educational materials

My authored and co-authored books primarily serve students in technical faculties, aiming to facilitate the acquisition of new knowledge, deepen understanding of complex theoretical concepts, and develop practical skills. These works cover a wide range of essential domains including theoretical computer science, advanced programming paradigms, computational algorithms, digital image processing, and numerical analysis methods. The published works include:

1. Elemente de C++ modern. Ghid practic pentru curs și laborator (2025), **A. Băicoianu**, A. Tamaş, A. Kerestely, I. Popa, V. Vrabie, ISBN: 978-606-19-1654-2
   A comprehensive guide focusing on modern C++ programming concepts, designed to support both theoretical understanding and practical implementation. The book addresses the transition from traditional to modern C++ programming paradigms.

2. Procesarea imaginilor digitale (2024), I.C. Plajer, C.G. Fieraru, **A. Băicoianu**, ISBN: 978-606-19-1720-4
   A comprehensive guide focusing on digital image analysis and processing, particularly relevant for understanding artificial intelligence aspects, especially convolutional networks and image filtering techniques.

3. Algoritmi fundamentali. Curs pentru învățământul la distanță (2023), **A. Băicoianu**, L. Majercsik, I.C. Plajer, ISBN: 978-606-19-1654-2
   This ebook presents fundamental algorithms through a carefully selected series of C++ examples, systematically organized by level of difficulty. The examples are designed to illustrate core algorithmic concepts, while promoting structured thinking and efficient coding practices.

4. Algoritmi fundamentali – Ghid practic pentru curs si laborator (2021), **A. Băicoianu**, L. Majercsik, ISBN: 978-606-19-1347-3
   A comprehensive guide to fundamental programming methods and algorithmic techniques, tailored for students in Mathematics and Computer Science. The material combines theoretical concepts with practical C++ implementations, offering a gradual learning path through examples of increasing complexity.

5. Metoda elementului finit. Implementare cu FreeFem++. Îndrumar de laborator. (2017), **A. Băicoianu**, ISBN: 978-606-19-0895-0
   Laboratory guide connecting partial differential equations with practical applications using FreeFem++.

6. Limbaje formale și teoria automatelor. Culegere de probleme (2012), **A. Băicoianu**, I.C. Plajer, ISBN: 978-606-19-0076-3
   Exercise collection supporting concepts in modeling, formal languages, and compilation techniques.

7. Limbaje formale şi teoria automatelor (2013), D. Marinescu, **A. Băicoianu**, I.C. Plajer
   Distance learning course material with concrete examples and detailed explanations.

8. Tehnici de compilare (2013), D. Marinescu, **A. Băicoianu**, I.C. Plajer
   Distance learning course material for compilation techniques.

With these projects, I want to help the next generation of professionals in computer science and engineering with both foundational knowledge and practical expertise.

Additional professional activities and contributions

Beyond research and teaching, I am involved in various technical and educational initiatives that combine academic expertise with practical applications. These activities include software development, knowledge dissemination, specialized training, and participation in international projects, as detailed below:

Category I - Software development and technical contributions

- Contributor to MatPopMod (version 0.2), a state-of-the-art Python library for matrix population models studies (https://bienvenu.gitlab.io/matpopmod/index.html). This collaboration was developed with original authors Dr. François Bienvenu (former postdoc at *Transilvania University*, currently researcher at Université de Bourgogne Franche-Comté, Besançon) and Dr. Guilhem Doulcier (postdoc at Macquarie University, Australia). The contribution was enhanced through a research visit to ETH Institute, Zurich (August 11-17, 2023, https://eth-its.ethz.ch/its-guests0/2023.html?cq_ck=1673368206541), and continues with active participation in the library's ongoing development.

- Co-author of specialized book chapter for farmers in "Cartoful în România" (Vol. 33, 2024),
  https://potato.ro/_publicatii_files/cartoful_in_ro/2024CartofulInRomania_vol.33%20.pdf

Category II - Keynote speaking and training events

- Speaker, DITARTIS Project Training Event, 16-17 January 2025, Universitat Politècnica de València
  Presentation: "Beyond the Black Box: xAI and Transfer Learning for Smarter Bearing Fault Diagnosis"
  A collaborative research effort with Universitat Politècnica de València partners Jose

E. Ruiz Sarrió and Jose Alfonso Antonino-Daviu, integrating *Siemens Industry Software*'s machine learning expertise with industrial engineering knowledge in condition monitoring and bearing fault diagnosis. The work represents a synergistic approach combining advanced ML methodologies with domain-specific industrial engineering applications.

- Speaker, AI4AGRI Summer School, 14–19 July 2025
  Presentations: "DACIA5 Data Set and Applications" & "ResNet18-based Crop Identification" [79]
  Notable outcome: Initiated the DACIA5 Data Challenge, an open research competition designed for students and early-career researchers in the fields of machine learning, Earth observation, and agricultural applications. The challenge promotes innovative utilization of the DACIA5 dataset through model development, data fusion approaches, and exploration of novel research directions. Also, an extended version of this challenge will be featured as a task on agriculture and images in ImageCLEF 2026.

- Speaker, SEEN Project Pilot Course, R&D Institute of *Transilvania University* of Braşov
  The presentation "AI Tools for Entrepreneurs" [104] was part of a comprehensive entrepreneurship training program that covered AI applications, research project management, and team leadership skills for both students and professors.

- Speaker, Workshop: NDVI Mapping for Agricultural Crop Monitoring at National Institute of Research and Development for Potato and Sugar Beet (INCDCSZ) Braşov, Romania, February 6, 2025. Delivered a specialized workshop focusing on NDVI mapping techniques for agricultural applications. This collaborative workshop facilitated knowledge transfer between academic research and practical agricultural applications, strengthening institutional partnerships in precision agriculture.

- Speaker, Workshop: Anomaly and specificity in agricultural crops at National Institute of Research and Development for Potato and Sugar Beet (INCDCSZ) Braşov, Romania, February 23, 2024. The specific session "Exploring anomalies in satellite imagery within the context of precision agriculture" focused on methods for identifying and interpreting anomalies in satellite imagery, which is important for optimizing different tasks in precision agriculture.

- Member of *Knowledge Committee* for Cyber Agent Project, Project ID: 101111732, Romanian coordinator: Prof. Toma Cristina.

These activities demonstrate my commitment to advancing both technical and educational initiatives across multiple domains, particularly in the integration of AI technologies with practical applications and entrepreneurial development.

Chapter 1. AI in Earth Observation and agricultural analytics

Problem statement and research context

AI and ML are rapidly transforming the way agricultural and Earth observation data are acquired, processed, and interpreted. These technologies enable a new paradigm of intelligent analysis in remote sensing, offering tools to monitor crops, detect anomalies, and optimize decision-making with unprecedented precision.

This chapter presents a set of research contributions developed within the framework of the AI4AGRI research coordinated by *Transilvania University* of Brașov. The project is grounded in the recognition that Romania has significant agricultural potential and could play a key role in Europe's strategic autonomy, food security, and sustainability. However, this potential remains underexploited due to several systemic issues, including limited access to technology and funding, soil fragmentation and erosion, and the ongoing desertification of certain regions. In response, AI4AGRI utilizes the power of AI-based systems and the availability of Copernicus Earth Observation (EO) data to build modern solutions for agriculture. Funded by the European Union, the project aims to establish a dedicated research center for AI on EO data in the agricultural sector, through twinning with leading institutes in artificial intelligence and remote sensing from France and Italy. Beyond its scientific objectives, the center is designed to provide training for young researchers, generate vegetation status maps for farmers, and enhance research and innovation management capabilities in the region. Within this strategic framework, this chapter highlights several key scientific outcomes that align with the mission and objectives of AI4AGRI. These contributions have emerged from a sustained research effort and address the main stages of intelligent agricultural analysis based on satellite data. The content follows a logical progression:

1. It begins with the development of intelligent visualization methodologies for spectral image interpretation and anomaly detection.

2. Then it explores methods for aggregating and harmonizing multisource data (Sentinel-1, Sentinel-2, and ground-based information). This is followed by the construction of a specialized labeled dataset (DACIA5) to support crop classification.

3. Finally, it investigates model adaptation and data augmentation techniques in deep learning pipelines, evaluating their crop-specific impact on classification accuracy and NDVI-based analysis.

Together, these works contribute to the advancement of AI-enhanced agricultural monitoring and represent a coherent scientific response to the broader socio-economic and technological challenges addressed by the AI4AGRI project.

I. Intelligent visualization and vegetation index-based interpretation

Recent advances in AI have significantly improved the way remote sensing data is visualized, enabling more interpretable and actionable outputs. This section focuses on intelligent visualization methods applied to multispectral and hyperspectral satellite imagery, emphasizing the use of vegetation indices - particularly the *Normalized Difference Vegetation Index* (NDVI) - as part of the processing pipeline.

The included studies explore both classical and ML-based approaches for multisource data visualization and spectral interpretation. An emphasis is placed on how dataset characteristics (such as variability and quantity) influence visual output quality and on how color calibration techniques, like Macbeth charts, enhance consistency. Furthermore, efficient NDVI-based methods are proposed to support automated anomaly detection and early crop health monitoring. Comparative evaluations between NDVI outputs from different satellite missions, including Sentinel-2 and *PRISMA*, are also presented.

To support these directions, this section includes a selection of research papers that apply and evaluate the discussed methods. Each article is briefly discussed below, highlighting the main findings, methodological innovations, and practical implications in a structured format, with more detailed analysis provided in the subsequent subsections.

Highlighted studies and key findings:

- I.C. Plajer, **A. Băicoianu**, L. Majercsik, M. Ivanovici, 2024 - *Multisource Remote Sensing Data Visualization using Machine Learning* [133]

  Main findings:

  - Demonstrates effective processing of aggregated multispectral datasets from diverse sources using a fully connected neural network (FCNN), with a strong focus on the impact of the preprocessing stage.
  - A secondary objective addresses the need for consistent and accurate rendering of spectral images into the RGB color space, aiming to enhance interpretability and support visual inspection.
  - The FCNN model achieved near state-of-the-art performance - ranking just slightly below the multi-path convolutional neural network (MPCNN) in quantitative metrics - while offering noticeably superior visual realism in the generated outputs.
  - The best visual contrast and image quality were obtained from FCNNs trained on the *CAVE* dataset, particularly when adaptive preprocessing was applied during inference.

- I.C. Plajer, **A. Băicoianu**, L. Majercsik, 2024 - *AI-Based Visualization of Remotely-Sensed Spectral Images* [131]

  Main findings:

  - Proposes a neural network–based approach for the visualization of multispectral and hyperspectral satellite images using a FCNN. The FCNN was trained on two publicly available datasets, ensuring generalization and robustness across varying spectral inputs.
  - The method enables consistent and visually realistic mapping of spectral data into the RGB domain, addressing a key limitation in human- and display-readable satellite imagery. Experimental results showed that the FCNN approach outperforms traditional visualization techniques in terms of visual quality and end-to-end performance.
  - The successfully application on *PRISMA* hyperspectral data, highlights its applicability in real-world scenarios, particularly for agricultural and environmental monitoring.

- I.C. Plajer, **A. Băicoianu**, M. Ivanovici, 2025 - *Multispectral Image Visualization and NDVI-based Anomaly Detection* [130]

  Main findings:

  - Demonstrates that the size and variability of training datasets have a greater impact on the quality of multispectral image visualization than the accuracy of labels or data quality.
  - Highlights the use of Macbeth color charts for improving dataset consistency and color calibration in spectral analysis.
  - Proposes a simple and computationally efficient methodology for anomaly detection in multispectral images, based on the NDVI.
  - Shows that the NDVI-based anomaly detection can automatically identify potential issues in crop development, supporting timely agricultural health monitoring and intervention.
  - Provides practical insights into how dataset characteristics influence visualization performance and proposes solutions to enhance satellite imagery analysis for precision agriculture.

- I.C. Plajer, **A. Băicoianu**, L. Majercsik, M. Ivanovici, 2023 - *NDVI Computation from Hyperspectral Images* [132]

  Main findings:

  - The NDVI calculated from hyperspectral images demonstrates robustness and consistency across various methods, including selecting a single band near 660 nm for RED and near 830 nm for NIR, averaging spectral bands within the RED [645.9–684.1 nm] and NIR [785.6–892.1 nm] ranges, as well as applying Gaussian-weighted averaging across these spectral intervals.

- A comparative analysis of NDVI values derived from spectral imagery obtained from different sensor sources is presented.
- The study concludes that variations in sensor type and product processing levels significantly affect NDVI computations and consequently influence the interpretation of crop health. Specifically, the *PRISMA* L2C product tends to yield a more optimistic NDVI estimation, whereas the Sentinel-2 L2A product reflects a comparatively conservative assessment.

I.1. Technical aspects of the intelligent visualization pipeline

Introduction in hyperspectral and multispectral imaging

Multispectral (MS) and hyperspectral (HS) imaging techniques capture detailed spectral information about the Earth's surface by recording reflected electromagnetic radiation across multiple wavelengths. Unlike standard RGB images from conventional digital cameras, which contain only visible light information, images produced by MS and HS sensors include additional spectral data that contains information about the chemical composition of objects. This additional information makes these images extremely useful in a series of applications in areas such as agriculture [181], [3], forestry [145], environmental monitoring and ecology [48], object detection [82], land cover classification [37], [112], as well as military and industrial fields. The key distinction between MS and HS images lies in their spectral resolution. HS images have far higher spectral resolution, capturing data in several narrow and contiguous bands, whereas MS images only capture data in a small number of discrete bands. This higher spectral resolution allows for a more detailed characterization of surface materials based on their spectral signatures. MS and HS images offer valuable insight into the Earth's surface, and the choice between them depends on the specific application requirements and the level of spectral and spatial detail needed for analysis.

The diversity of sensor types employed in MS and HS imaging introduces specific technical challenges and data variations. Sensor configurations exhibit considerable variation in spectral band arrangement, wavelength ranges, and bandwidth specifications. This heterogeneity in data sources and sensor characteristics can result in significant inconsistencies when attempting to compare or integrate data from multiple sensors. Consequently, careful consideration of these variations is essential to ensure accurate and meaningful analysis outcomes.

Sensor calibration represents a critical aspect in both MS and HS imaging systems, directly impacting data accuracy and reliability. Calibration methodologies demonstrate significant variation across different sensors and platforms. These disparities in calibration approaches can introduce variations in radiometric and spectral accuracy of the acquired data. Implementation of appropriate calibration and validation protocols is therefore essential for establishing reliable and consistent datasets. Individual sensors possess unique operational characteristics and inherent limitations. These include specific spec-

tral response functions, signal-to-noise ratios, and dynamic ranges. Such characteristics significantly influence the quality and reliability of the captured data. Furthermore, spatial resolution varies among spectral imaging sensors, affecting the level of detail preserved in the acquired images. The integration or comparison of multisource sensor data with varying spatial resolutions presents additional challenges, particularly when spatial variability exhibits poor alignment.

The visualization of multispectral and hyperspectral data represents a fundamental first step in the interpretation process. These imaging systems capture spectral information that extends well beyond what standard RGB displays can show. While conventional displays are limited to three color channels, MS and HS data contain valuable information across multiple wavelengths, requiring specialized approaches for effective visualization and interpretation.

Several visualization techniques have been proposed in the scientific literature in order to generate realistic RGB images from spectral images. These include band selection [95], [164], independent component analysis (ICA) [197] and principal component analysis (PCA) [169], [87] based methods, linear and nonlinear methods [98], and, relatively recent, machine learning approaches [49], [168]. Accurate visualization of satellite images helps experts to directly spot problems in the field. For example, in agriculture, well-displayed images make it easier to identify areas where crops might be diseased or stressed. When satellite images are shown in a way that looks natural and realistic, they also help verify if more sophisticated analysis methods are working correctly. This is particularly useful for checking classification results - researchers can quickly see if the computer's analysis matches what they can see with their own eyes before diving into more detailed technical evaluations.

The study [133], co-authored by the author of the present thesis, introduced a methodology for processing multispectral and hyperspectral imagery from diverse sources. This work specifically targets the challenges of handling data variability across different sensors through a fully connected neural network approach. Although data visualization was implemented as a secondary objective, it proved valuable in validating the fusion methodology. This methodology focuses on developing effective normalization and standardization techniques for diverse spectral datasets. The proposed FCNN approach demonstrates versatility in processing both multispectral and hyperspectral imagery, while effectively addressing band selection and spectral variability challenges. The network architecture and training parameters have been carefully optimized for these purposes.

The main advantages of the proposed method lie in the generalization of the approach and adaptability to data of different spectral and spatial resolution, allowing to effectively handle both MS and HS images. Furthermore, by combining publicly available datasets with appropriately labeled data for network training, a principled and accessible foundation for model learning is established. The proposed method also has the advantage of not requiring further image post-processing of the obtained RGB image, which

leads to a more polished and refined output straight out of the model.

Datasets

The study [133], co-authored by the author of the present thesis, utilized multiple datasets for both training and testing purposes. For the training phase, two well-established datasets were selected: *CAVE* and *UGR*. The *CAVE* dataset [192] comprises 32 multispectral images captured in controlled indoor conditions, with a spatial resolution of 512×512 pixels and 31 spectral bands spanning 400-700 nm. The *UGR* dataset [52] contains 14 outdoor urban scenes at 1000×900 pixels resolution, covering a spectral range of 400-1000 nm across 61 bands, with 31 bands in the visible spectrum. Both datasets provide corresponding RGB images, making them particularly suitable for this analysis.

Testing and validation were conducted using three distinct data sources. Two hyperspectral images were acquired from the *PRISMA* satellite, captured over Braşov county, Romania, in October 2022 and March 2023. These images feature 239 spectral bands ranging from 400 nm to 2500 nm, with a spatial resolution of 1000×1000 pixels and 30 m ground sampling distance [5]. The *Pavia University* dataset, acquired by..., was also employed in the study. Additionally, the *Pavia University* dataset, acquired by the ROSIS sensor, was utilized. This dataset provides 103 usable spectral bands between 430–860 nm, with a spatial resolution of 1.3 m across a 610×610 pixel area [60].

In order to test the *PRISMA* and the *Pavia University* images on the FCNN trained on the *CAVE* or the *UGR* dataset it is necessary to adapt these images to the network's input layer. As this input is calibrated to receive pixels with the spectral signature of *CAVE* images, each spectral pixel of the test image has to be mapped to the wavelengths of a *CAVE* image pixel. This has been done in this study by linear interpolation. For each wavelength of a test image (*PRISMA* or *Pavia University*), the value of a *CAVE* image channel is interpolated from the values of the two neighboring channels of the test image.

Fig. 1 shows an original HS pixel of a *PRISMA* image together with the interpolated values for this pixel. It can be seen in Fig. 1, that by linearly interpolating the *PRISMA* HS image to fit the bands of the *CAVE* dataset, the changes in the data profile are negligible, thus justifying this approach.

Methodology, experimental results, and comparative quality assessment

In the study [133], in which the author of the present thesis is a co-author, multiple data preprocessing and transformation approaches were explored, with particular emphasis on two strategies that demonstrated the most promising results. One of the strategies involved preprocessing of each file, concatenation of data from all the input files, shuffling, and separation into train subset and test subset (respectively $\frac{3}{4}$ and $\frac{1}{4}$ of all existing pixels). In the following, this strategy is referred to as *individual preprocessing*. The second approach, termed *global preprocessing*, involves concatenating the raw

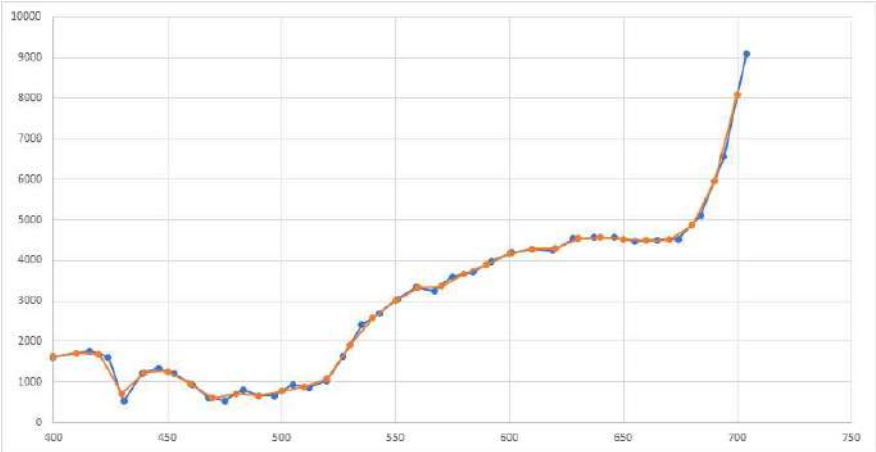

Figure 1: Original values of one HS pixel from the first *PRISMA* image (in blue color) and interpolated values (in orange color).

data without initial preprocessing, followed by shuffling and train-test splitting (maintaining identical proportions as the previous strategy). Preprocessing methods (standardization or normalization) are then applied exclusively to the training subset. During testing, the same preprocessing parameters from the training phase are applied to the test data, ensuring consistent data scaling across the entire pipeline. From a technical perspective, both strategies, *StandardScaler()* and *MinMaxScaler()*, were applied for standardization and normalization, respectively.

A fully connected neural network was selected to address the regression problem of mapping n-dimensional spectral pixels to RGB space. As shown in Fig. 2, the network architecture comprises five layers: an input layer, three hidden layers, and an output layer.

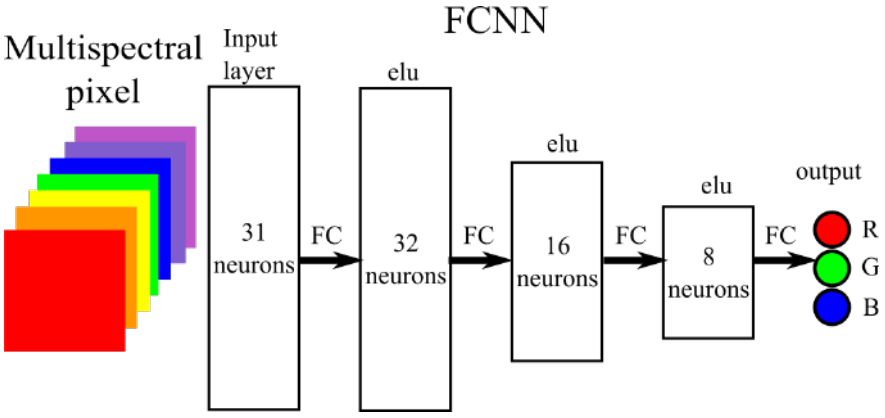

Figure 2: The network structure.

The network consists of a 31-neuron input layer, matching the number of visual spectral bands in both *CAVE* and *UGR* datasets, and a 3-neuron output layer corresponding to RGB channels. The hidden layers follow a decreasing pattern (32, 16, and 8 neurons) using ELU activation functions to preserve negative values.

The Huber loss function was selected as a performance metric, as it provides a balanced approach between *Mean Squared Error (MSE)* and *Mean Absolute Error (MAE)*, effectively handling outliers while maintaining sensitivity to smaller errors. The Huber loss

is defined as follows in Equation 1.

$$L_\delta(y, f(x)) = \begin{cases} \frac{1}{2}(y - f(x))^2, & |y - f(x)| \leq \delta \\ \delta(|y - f(x)| - \frac{1}{2}\delta), & |y - f(x)| > \delta \end{cases} \quad (1)$$

where y is the true value, $f(x)$ is the predicted value, and δ is a parameter that determines the threshold for which, in the loss function, the MSE is replaced by MAE. The Huber loss function combines the advantages of both MSE and MAE through an adaptive threshold mechanism. In the current implementation, δ was empirically set to 10.0, aligning with human perception principles similar to the Just Noticeable Difference concept in color science and corresponding to $\Delta E = 3$ in CIELab [133].

The model was trained for 150 epochs with a batch size of 2048 pixels, using Adam optimizer with learning rate $\alpha = 0.005$. Training was performed on an Intel i7-7700 CPU, taking approximately 2 hours for the *UGR* dataset. Data preprocessing followed the *global standardization* approach described before.

In order to enhance clarity and facilitate a deeper understanding of the proposed methodology, pseudocode has been provided to describe key steps in detail in Algorithm 1.

As shown in Fig. 3, the loss function exhibits rapid initial decay on the training set, followed by stabilization, while the test set requires additional epochs to achieve comparable convergence for both *CAVE* and *UGR* datasets.

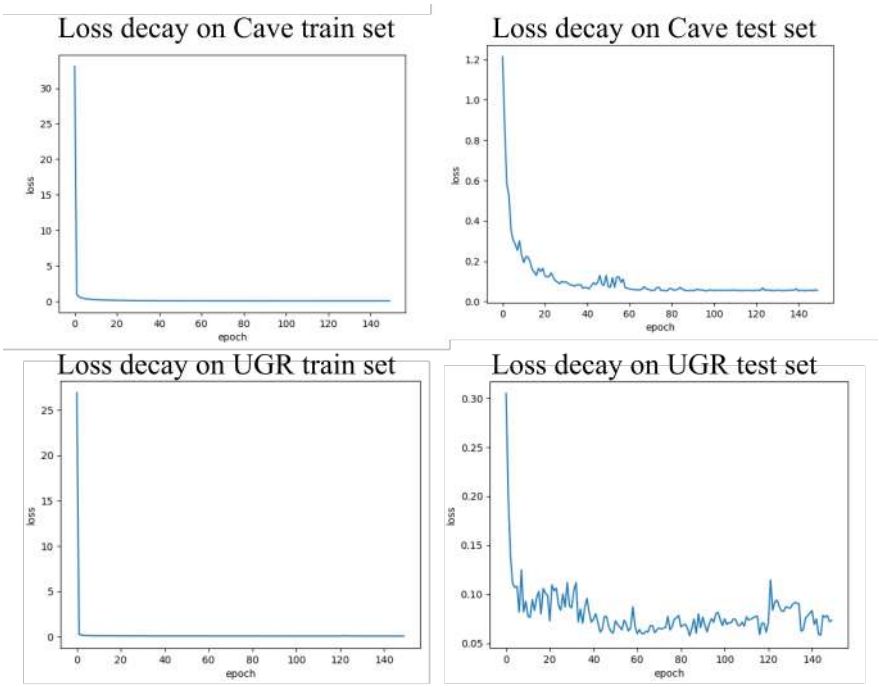

Figure 3: The loss decay on the train and test sets when training the model with the *CAVE* dataset, respectively *UGR* dataset.

In the study [133], co-authored by the author of the present thesis, two global preprocessing approaches were explored: one with normalization and one with standardization.

Algorithm 1 Pseudocode for the training of the FCNN.

Input: $path\_dataset,\ model,\ lr = 0.005,\ epochs = 150$
$(data, labels) \leftarrow load\_pixels(path\_dataset)$
$shuffle(data, labels)$
$nr\_train\_data \leftarrow (\frac{3}{4}) * len(data)$
$nr\_test\_data \leftarrow len(data) - nr\_train\_data$
$train\_set \leftarrow (data.head(nr\_train\_data), labels.head(nr\_train\_data))$
$test\_set \leftarrow (data.tail(nr\_test\_data), labels.tail(nr\_test\_data))$
Initialize the StandardScaler and Standardize training data
Apply Scaler on test data
$batch\_size \leftarrow 2048$
Initialize weights of the $model$
$optimizer \leftarrow AdamOptimizer(lr, weight\_decay = 0.0008)$
$best\_loss \leftarrow maxValue$
$loss\_function \leftarrow HuberLoss(10.0)$
$train\_loss\_vec \leftarrow []$
$test\_loss\_vec \leftarrow []$
for $epoch \in (1, epochs)$ **do**
    $train\_loss \leftarrow 0$
    for $batch, (x, y) \in train\_set$ **do**
        $y\_pred = model(x)$
        $loss \leftarrow loss\_function(y, y\_pred)$
        Backpropagation Step - backprop(loss)
        $train\_loss \leftarrow train\_loss + loss.item()$
    end for
    $train\_loss \leftarrow train\_loss/(batch\_size)$
    $train\_loss\_vec.append(train\_loss)$
    if $train\_loss < best\_loss$ **then**
        Save current weights
        $best\_loss \leftarrow train\_loss$
    else
    end if
    $test\_loss \leftarrow 0$
    for $batch, (x, y) \in test\_set$ **do**
        $y\_pred \leftarrow model(x)$
        $loss \leftarrow loss\_function(y, y\_pred)$
        $test\_loss \leftarrow test\_loss + loss.item()$
    end for
    $test\_loss \leftarrow test\_loss/(batch\_size)$
    $test\_loss\_vec.append(test\_loss)$
end for
Save current weights
$plot(train\_loss\_vec, test\_loss\_vec)$

Both methods were applied to the complete training set, with the resulting parameters subsequently used for test set and inference preprocessing.

Fig. 4 illustrates the impact of global preprocessing approaches on pixel color mapping performance. The visualization includes original samples from *CAVE* and *UGR* datasets (a), followed by FCNN mapping results using global normalization on *CAVE* (b) and *UGR* (c), and global standardization results on *CAVE* (d) and *UGR* (e) respectively.

Fig. 5 illustrates the impact of individual image preprocessing approaches on mapping performance. The visualization includes original samples from *CAVE* and *UGR* datasets (a), followed by FCNN mapping results using per-image normalization on *CAVE* (b) and*UGR* (c), and per-image standardization results on *CAVE* (d) and *UGR* (e) respectively.

For normalization-based preprocessing, the FCNN performance was evaluated on both *CAVE* (Fig. 4b) and *UGR* (Fig. 4c) sample images using min-max normalization. The model's generalization capabilities were tested on external datasets: the *Pavia University* image (Fig. 7a, Fig. 7b) and *PRISMA* satellite imagery (Fig. 6a, Fig. 6b). Similarly, for standardization-based preprocessing, the FCNN performance was analyzed on *CAVE*

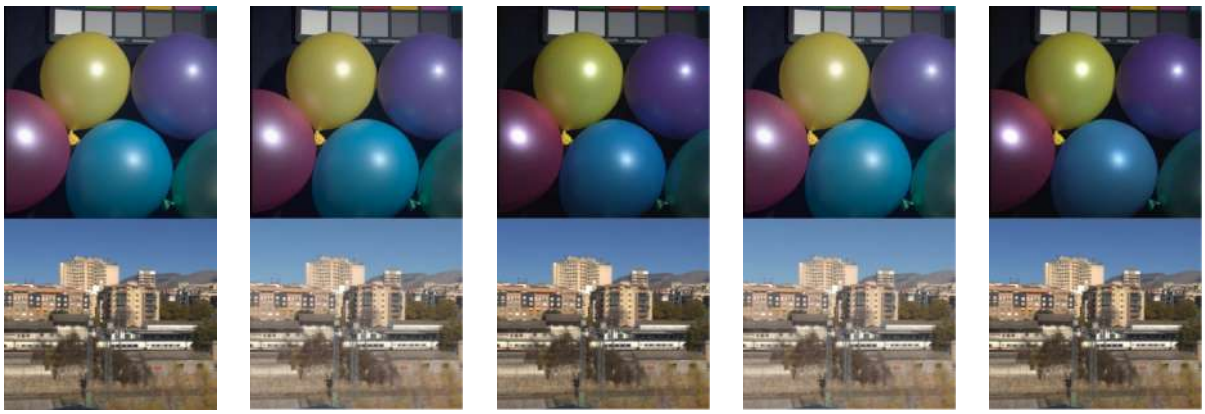

(a) (b) (c) (d) (e)

Figure 4: Output visualizations showing the effects of global preprocessing on FCNN mapping for the *CAVE* and *UGR* datasets.

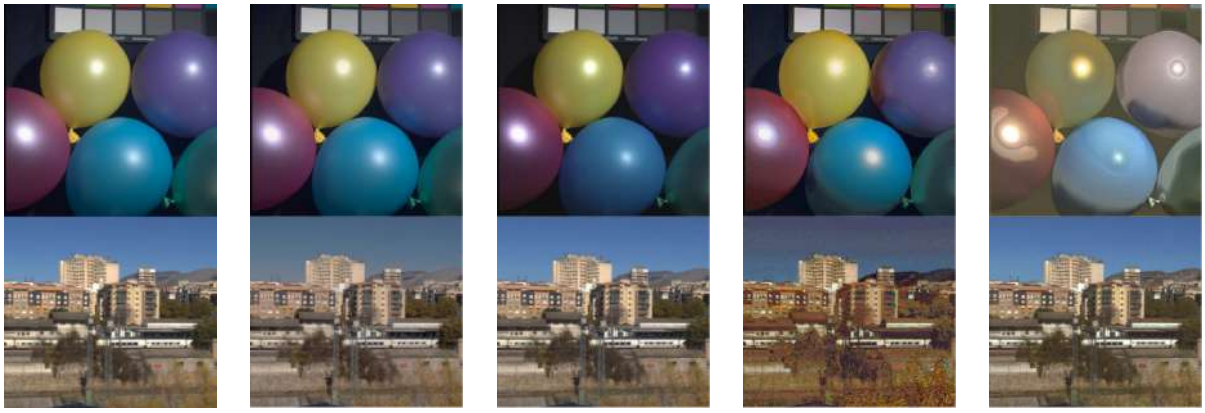

(a) (b) (c) (d) (e)

Figure 5: Output visualizations showing the effects of per-image preprocessing on FCNN mapping for the *CAVE* and *UGR* datasets.

(Fig. 4d) and *UGR* (Fig. 4e) samples. The model was then evaluated on *Pavia University* (Fig. 7c, Fig. 7d) and *PRISMA* data (Fig. 6c, Fig. 6d).

In the second approach, individual preprocessing strategies were implemented, with each image preprocessed independently before dataset concatenation and train-test splitting. Two methods were implemented: individual normalization, where each image was scaled to [0, 1] using its own min-max values, and individual standardization, where each image is transformed using its own mean and standard deviation. Results for both *CAVE* and *UGR* training scenarios are presented in Fig. 5. The methods were validated on external datasets, with results shown for *Pavia University* (Fig. 7e, Fig. 7f, Fig. 7g, Fig. 7h) and *PRISMA* imagery (Fig. 6e, Fig. 6f, Fig. 6g, Fig. 6h).

These experiments revealed that while global preprocessing (normalization or standardization) performed optimally on *CAVE* and *UGR* datasets, it yielded suboptimal results for *Pavia University* and *PRISMA* images. This discrepancy can be attributed to the distinct data distributions across datasets, as illustrated in Fig. 8, which compares the green channel (650 nm) distributions for the *CAVE* (Fig. 8a), *UGR* (Fig. 8b), *Pavia*

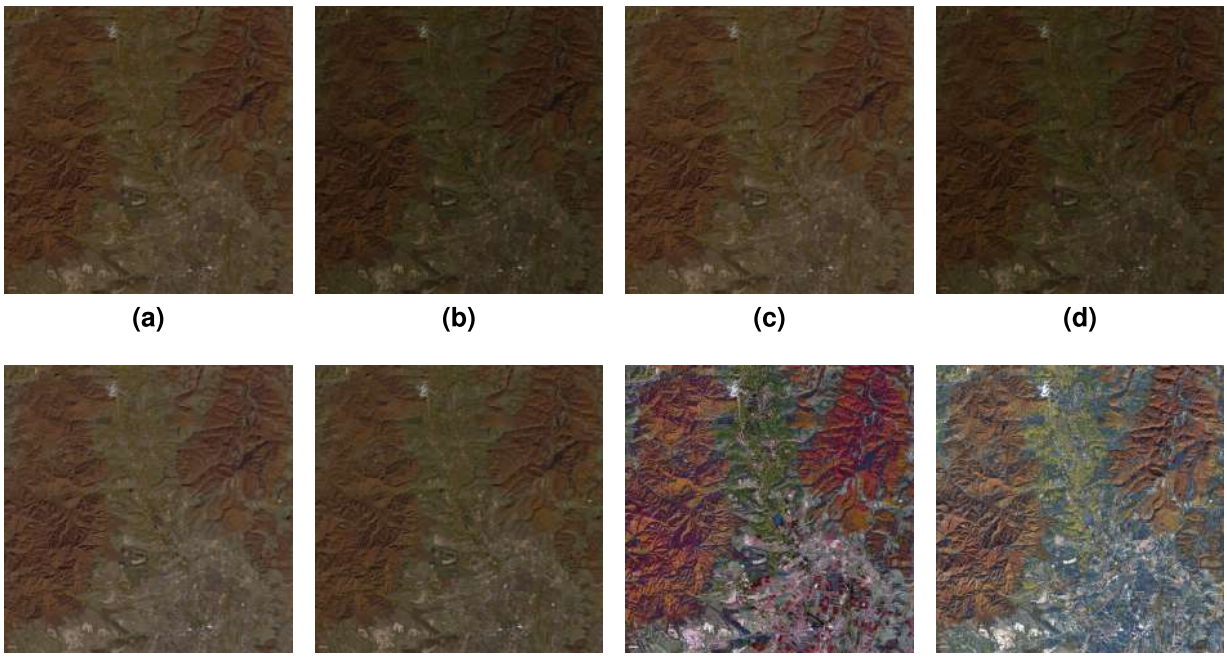

Figure 6: Output visualizations from the FCNN model applied to the first *PRISMA* HS image in the various preprocessing scenarios.

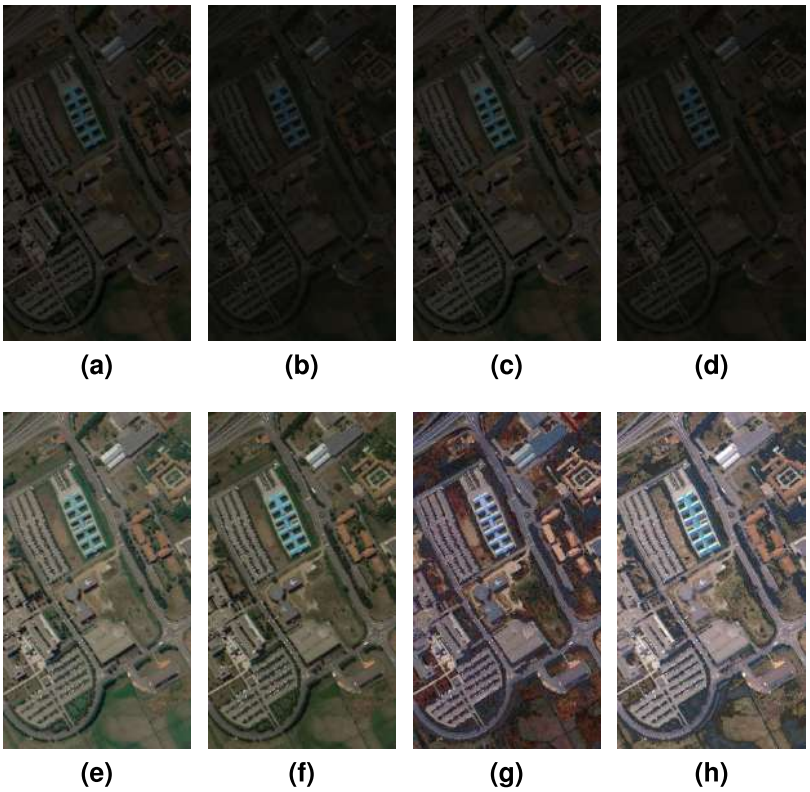

Figure 7: Output visualizations from the FCNN model applied to the *Pavia University* MS image in the various preprocessing scenarios.

University (Fig. 8c), and *PRISMA* (Fig. 8d) datasets. The green channel was selected for illustration, but the observation holds for any of the 31 other channels of the spectral images considered.

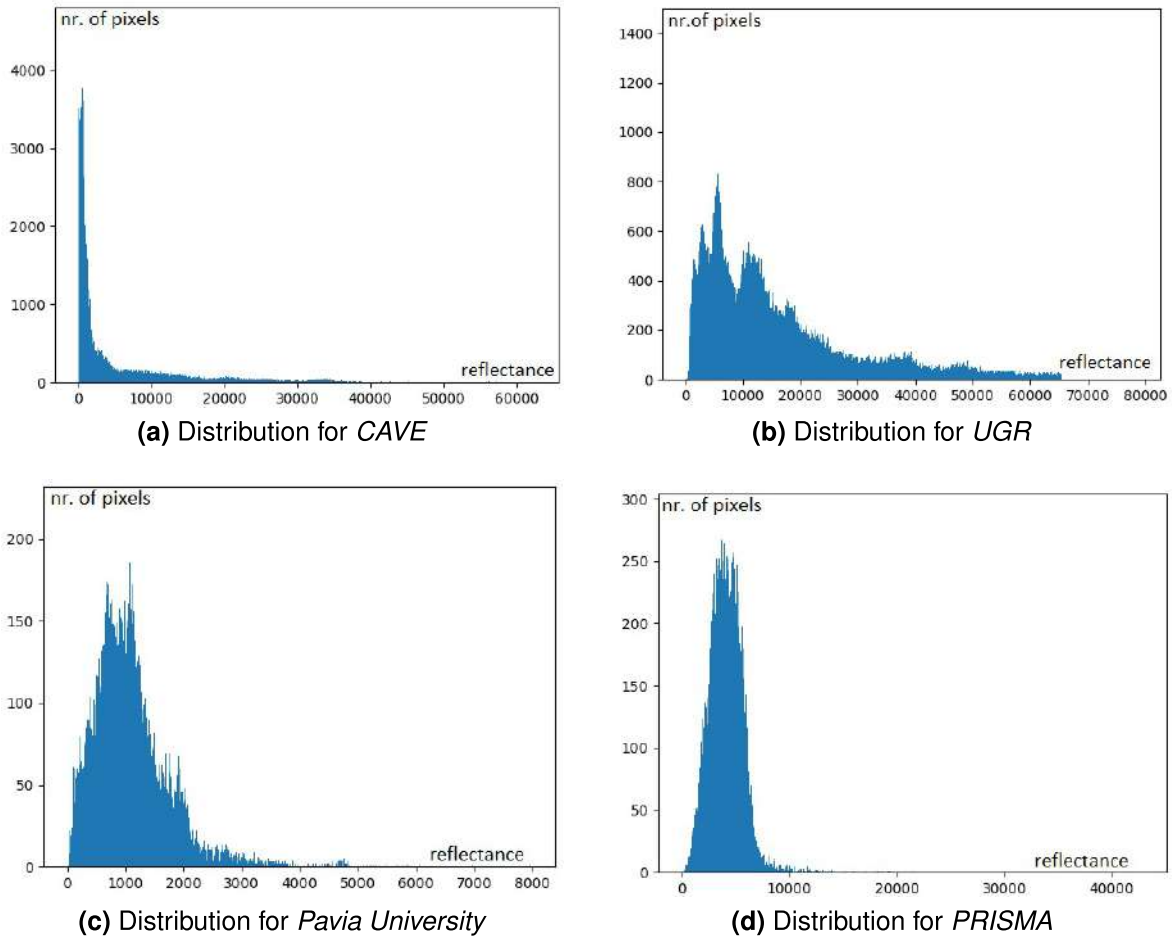

(a) Distribution for *CAVE*

(b) Distribution for *UGR*

(c) Distribution for *Pavia University*

(d) Distribution for *PRISMA*

Figure 8: Distributions on the green channel (considered at 550 nm) for *CAVE* and *UGR* datasets and for *Pavia University* and *PRISMA* spectral images.

To enhance model generalization across diverse data sources, an alternative approach was developed. While maintaining global standardization during training, image-specific standardization was applied during inference, with each new image standardized using its own statistical parameters. This ensures each channel is transformed to zero mean and unit standard deviation before processing through the trained FCNN model. For a detailed implementation overview, the inference procedure is formalized in Algorithm 2.

Algorithm 2 Pseudocode for visualization using the FCNN.

Input: $path\_ms\_image$, $model$, $input\_freq$, $target\_freq$
1: $image \leftarrow load(path\_ms\_image)$
2: $interpolated \leftarrow$ Interpolate$(image, input_freq, target_freq)$
3: Scale $interpolated$ with StandardScaler
4: $predicted \leftarrow model(interpolated)$
5: Make RGB image from $predicted$

The results of this approach are presented as follows: for the *Pavia University* image in Fig. 10a (*CAVE* trained FCNN) and Fig. 10d (*UGR* trained FCNN), for the first *PRISMA* image in Fig. 10b (*CAVE* trained FCNN) and Fig. 10e and for the second *PRISMA* image in Fig. 10c and respectively Fig. 12f.

Multiple approaches were evaluated using various spectral datasets from different

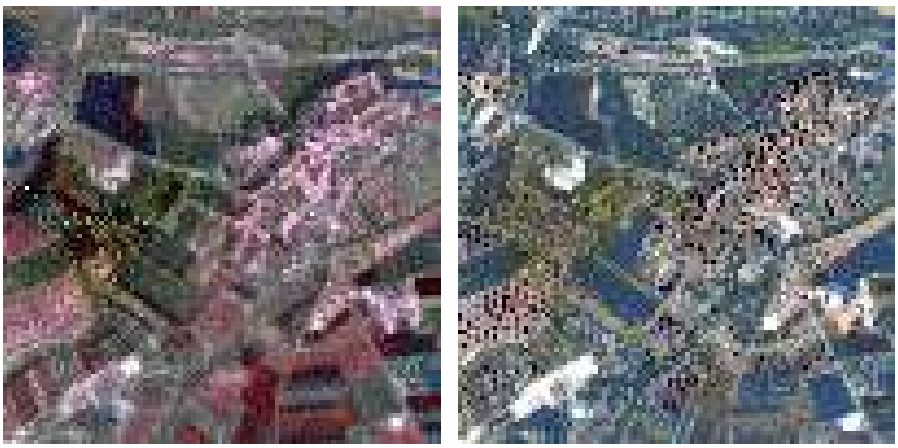

(a) Region from image in Fig. 6g **(b)** Region from image in Fig. 6h

Figure 9: Illustrations of the artifacts in the *PRISMA* image in a region of interest.

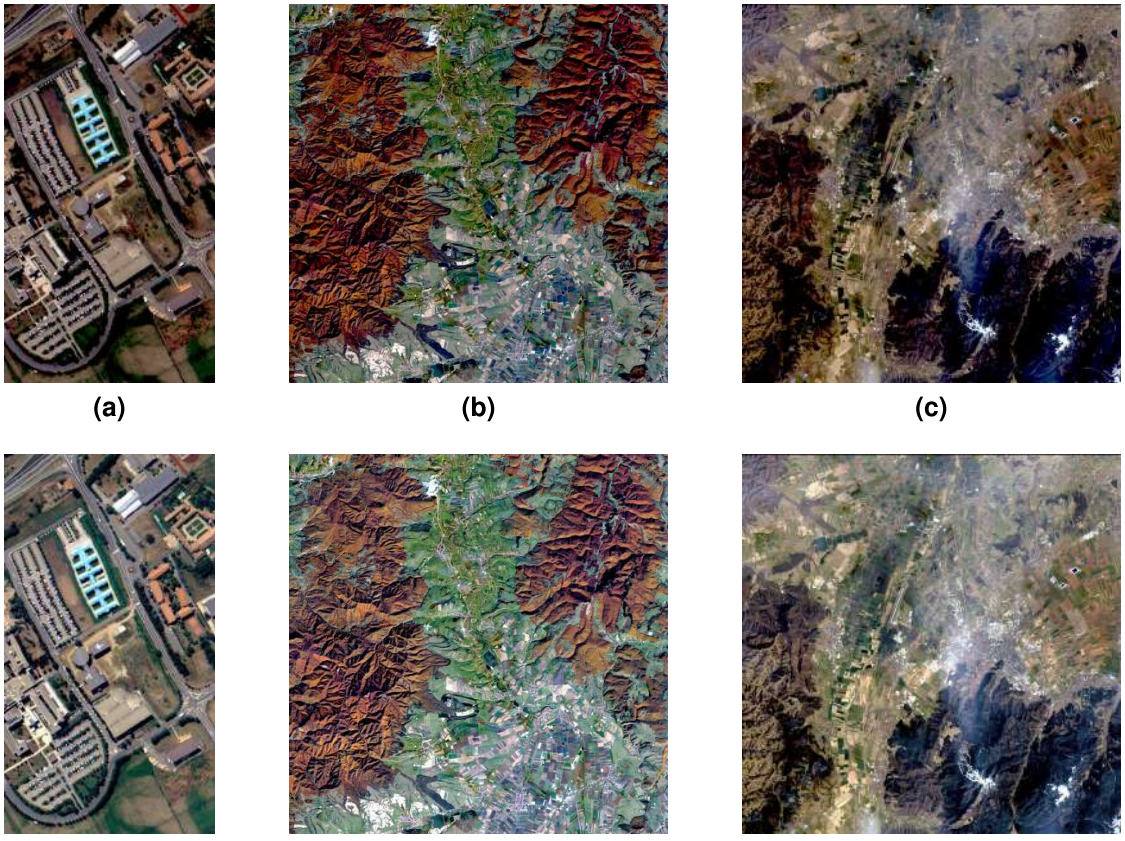

(a) **(b)** **(c)**

(d) **(e)** **(f)**

Figure 10: Visualization of the results on *Pavia University* and *PRISMA*, when standardizing those images according to their own distributions. Top: FCNN trained on *CAVE*; Bottom: FCNN trained on *UGR*..

sensors to determine the most effective visualization method. The experimental results (Fig. 4 to Fig. 13) demonstrate that global normalization during training produces consistent results, with better performance when training and test datasets match. Individual image normalization preserves feature relationships [7], while standardization can introduce artifacts (Fig. 9a, Fig. 9b).

For external datasets (*Pavia University* and *PRISMA*), which required spectral band interpolation, the best results were achieved using global preprocessing during training combined with image-specific standardization during inference (Fig. 10). The approach was compared with conventional methods, including band selection and XYZ space [105], decolorization-based [83], and MPCNN-based visualization [50] (Fig. 12, Fig. 13).

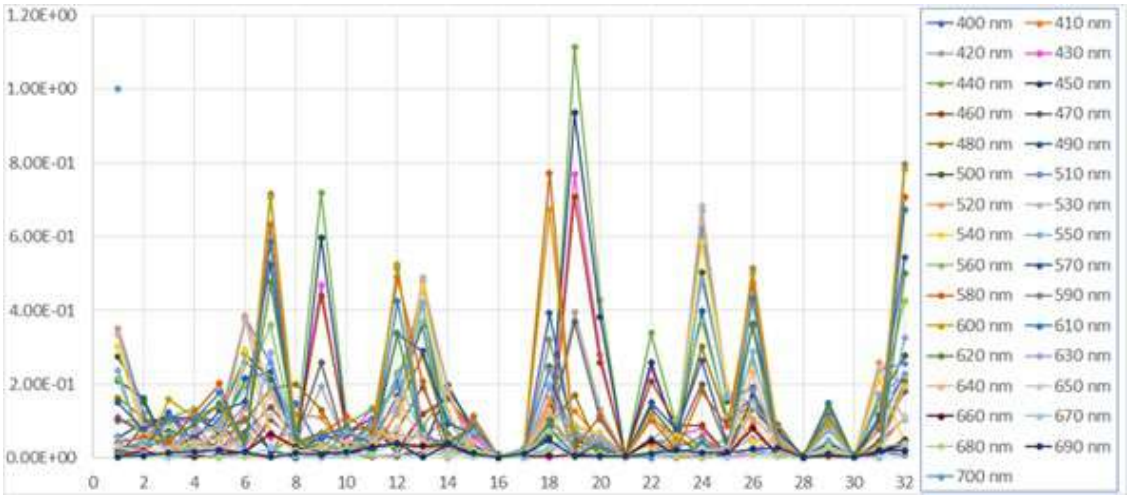

Figure 11: All the absolute values of the weights between the input and the first hidden layer.

The proposed method demonstrates superior performance compared to conventional approaches, especially for satellite imagery. By implementing image-specific standardization during inference, the proposed network effectively handles varying atmospheric conditions and input distributions, making it particularly valuable for satellite image visualization applications. Furthermore, it was investigated whether the approach, similar to band selection, favors certain wavelengths over others. All absolute values of the weights connecting the input to the first hidden layer were plotted, and the resulting visualization is presented in Fig. 11. It can be seen that this is not the case, demonstrating that each band contributes to the outcome. In Fig. 11, each color represents the weights of one input neuron to the neurons on the next layer.

In addition, in [133] the authors presented some numerical quality comparison between the results illustrated in Fig. 12 and Fig. 13, some no-reference image quality assessments [179, 184] like entropy, which estimates information quantity, fractal dimension, which estimates complexity and standard deviation, which estimates non-uniformity, were used. Figure 12 presents a comprehensive comparison of visualization techniques applied to the first *PRISMA* HS image. The methods range from traditional approaches such as band selection (a), XYZ space coloring (b), and decolorization-based visualization (c), to advanced techniques including MPCNN visualization (d) and the proposed FCNN approach trained on *CAVE* (e) and *UGR* (f) datasets. Figure 13 presents a comparative analysis of visualization techniques applied to the second *PRISMA* hyperspectral image. The comparison includes traditional methods such as band selection (a), XYZ space coloring (b), and decolorization-based visualization (c), alongside advanced techniques including MPCNN visualization (d) and the proposed FCNN approach trained on

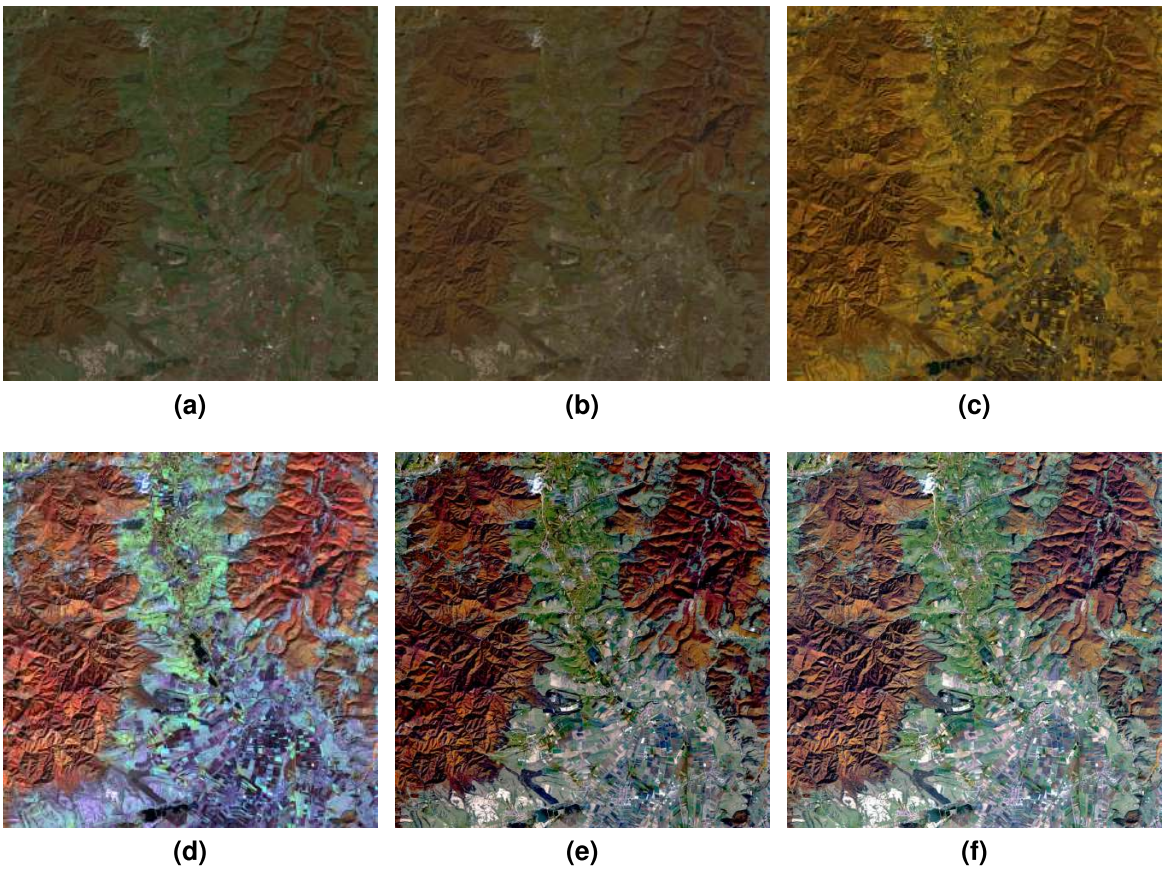

Figure 12: Visualization comparison of first PRISMA image: traditional methods (a-c), MPCNN (d), and proposed FCNN approach (e), (f).

Table 1: Comparative quality assessment for the first *PRISMA* image

Sample	*Entropy*	*Standard Deviation*	*Fractal Dimension with* L_{max}		
			41	71	101
FCNN / CAVE	15.6314	47.0056	3.7421	3.8271	3.8291
FCNN / UGR	15.8624	44.7105	3.7199	3.8064	3.7940
Band selection	12.6552	16.1236	2.7463	2.7484	2.8136
XYZ method	13.5262	20.4681	2.9795	2.9251	2.8930
Decolorization	14.0430	39.5901	2.9668	3.0568	3.1282
MPCNN	15.9284	48.0310	3.7245	3.9921	3.9566

CAVE (e) and *UGR* (f) datasets.

For calculating the fractal dimension, the parameters used were $LMAX \in \{41, 71, 101\}$, representing the maximum size of the hypercubes and a threshold of 0.00001 for the standard deviation [80]. Tables 1 and 2 present the quantitative evaluation of different visualization methods (Fig. 12 and Fig. 13). While MPCNN achieved highest entropy and variance values, it produced unnatural colors (Fig. 12d, Fig. 13d). The FCNN approach achieved comparable quantitative metrics while maintaining natural color representation, with *CAVE*-trained models showing slightly better contrast, validating the proposed methodology. These quantitative results further justify the approach proposed by the present approach.

The visualization of multisource spectral imagery remains challenging as satellite

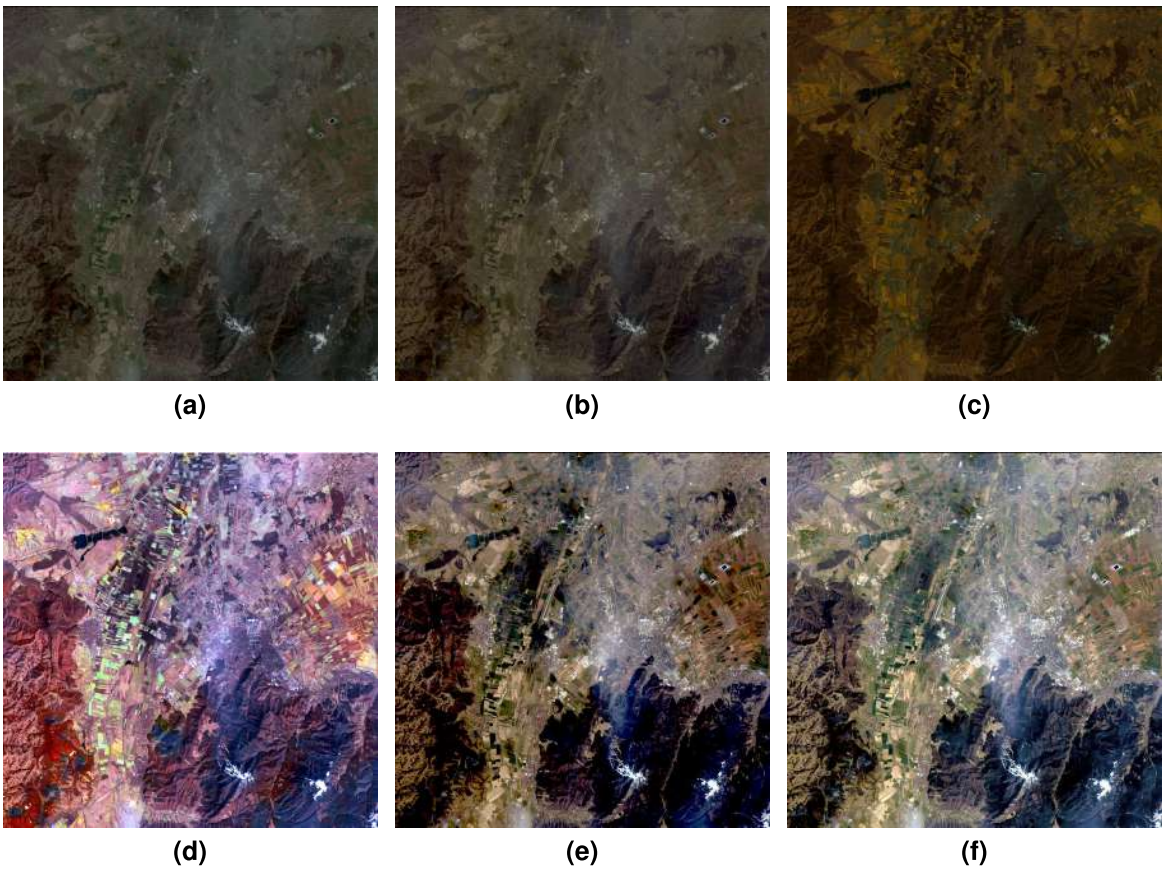

Figure 13: Visualization comparison of second PRISMA image: traditional methods (a-c), MPCNN (d), and proposed FCNN approach (e), (f).

Table 2: Comparative quality assessment for the second *PRISMA* image

Sample	Entropy	Standard Deviation	Fractal Dimension with L_{max}		
			41	71	101
FCNN / CAVE	14.9011	50.5844	3.1309	3.1925	3.2909
FCNN / UGR	15.5923	48.5822	3.0698	3.1678	3.2157
Band selection	14.2838	27.2990	2.4924	2.7755	3.1075
XYZ method	13.7575	20.7137	2.6343	2.7082	2.8104
Decolorization	12.9841	23.2493	2.6714	2.7646	2.8926
MPCNN	15.8832	53.4223	3.2908	3.4603	3.6736

data become increasingly accessible. In the study [133], coauthored by the present author of this thesis, this challenge was addressed through an FCNN-based approach trained on established public spectral datasets, with various preprocessing strategies implemented to handle diverse spectral signatures.

The method demonstrated strong performance both visually and quantitatively, effectively rendering agricultural and urban features. Although promising, the approach has its limitations: potential loss of color realism in low-variance images due to standardization, reliance on labeled training data, and the need for spectral band interpolation across different sources. Future work will focus on addressing these limitations and extending the methodology toward agricultural applications, particularly in vegetation index interpretation.

The complete results presented in this section (I.1), developed from initial work published in a Scopus-indexed conference paper [131], are part of an extensive study [133] conducted in collaboration with the AI4AGRI team, submitted to *IEEE Transactions on Geoscience and Remote Sensing*, a Q1-ranked journal in the field.

I.2. Technical aspects for spectral vegetation index computation

Recent advances in multi- and hyperspectral satellite imaging have revolutionized precision agriculture. Given the complexity of interpreting such data, vegetation indices provide valuable insights for agricultural applications. In this section, an approach [132] for computing vegetation indices from hyperspectral imagery is presented, comparing NDVI maps derived from Sentinel-2 and *PRISMA* data.

Problem statement and datasets

The analysis of vegetation health through spectral signatures has become crucial for precision agriculture [160], [159], enabling applications ranging from crop monitoring [152] to disease identification [100]. While vegetation indices like NDVI effectively transform complex spectral data into interpretable indicators, computing these indices from hyperspectral imagery presents unique challenges. This paper addresses the computation of NDVI from hyperspectral data, providing a comparative analysis with traditional multispectral-based NDVI calculations.

The study [132] utilizes both multispectral and hyperspectral imagery acquired on March 24, 2023. The multispectral data come from Sentinel-2, which captures 13 spectral bands and was resampled from 10 m to 30 m spatial resolution using SNAP [54]. The hyperspectral data are from *PRISMA*, which provides 239 spectral bands (400-2500 nm) at 30 m spatial resolution [5]. Three *PRISMA* product levels were analyzed: L1 (top-of-atmosphere radiance), L2B (surface-reflected radiance), and L2C (boundary reflection coefficient). RGB visualizations of both Sentinel-2 and *PRISMA* L2C products are presented in Fig. 14.

Normalized difference vegetation index (NDVI)

Vegetation indices are numerical values computed by mathematical combinations of the amplitudes of various wavelengths of light captured by remote sensing devices. They quantify different aspects of vegetation health, abundance, and photosynthetic activity. Since different aspects of plants, including their species characteristics, physiological conditions, and stage of growth, can influence the way different vegetation indices respond, a wide range of vegetation indices have been refined over time [23], with the objective of capturing as well as possible both intrinsic and extrinsic characteristics of plants.

The NDVI is one of the most widely recognized and commonly used vegetation indices. It takes advantage of the distinct reflectance characteristics of vegetation in the

(a) **(b)**

Figure 14: Sentinel-2 image (a) and *PRISMA* L2C image (b).

near-infrared (NIR) and red (RED) spectral bands, and is calculated using following Equation 2.

$$NDVI = \frac{NIR - RED}{NIR + RED} \tag{2}$$

NDVI values range from -1 to +1, where higher positive values typically indicate healthier, denser vegetation. While NDVI computation is straightforward for multispectral data with discrete bands, hyperspectral sensors present unique challenges and opportunities. Their higher spectral resolution offers multiple bands within the RED and NIR regions, requiring careful consideration of band selection or averaging strategies.

The computation of NDVI from multiple sensors requires careful consideration of their varying spectral and spatial resolutions, as well as atmospheric effects. To address these challenges with the hyperspectral data, which provides multiple bands in the RED and NIR regions, three computational approaches were developed:

(i) Band selection in the *PRISMA* images of the closest bands to the RED and NIR central wavelengths of sentinel B4 and B8. These are for *PRISMA* 664.8 nm for RED and 838.5 nm for NIR.

(ii) An average of the RED and NIR bands of *PRISMA* images over the intervals [645.9, 684.1] nm and [785.6, 892.1] nm respectively, similar to the spectral resolutions of the B4 and B8 bands for Sentinel-2 [53].

(iii) A weighted average using a Gaussian function over the RED and NIR bands of *PRISMA* images in the same intervals as in the unweighted average calculation.

For comparison, NDVI was computed also on the Sentinel-2 image by selecting B8 for NIR and B4 for RED bands [94].

In remote sensing, color palette selection significantly influences data interpretation and understanding. For NDVI visualization in agricultural applications, palettes commonly

transition from red (unhealthy) to green (healthy) vegetation, with variations ranging from simple binary schemes to complex gradients [127]. The proposed analysis utilizes the standardized Sentinel color palette [151] (Fig. 15a) for consistent NDVI representation. This refers specifically to the palette version in use at the time of experimentation, as more recent updates may include modifications.

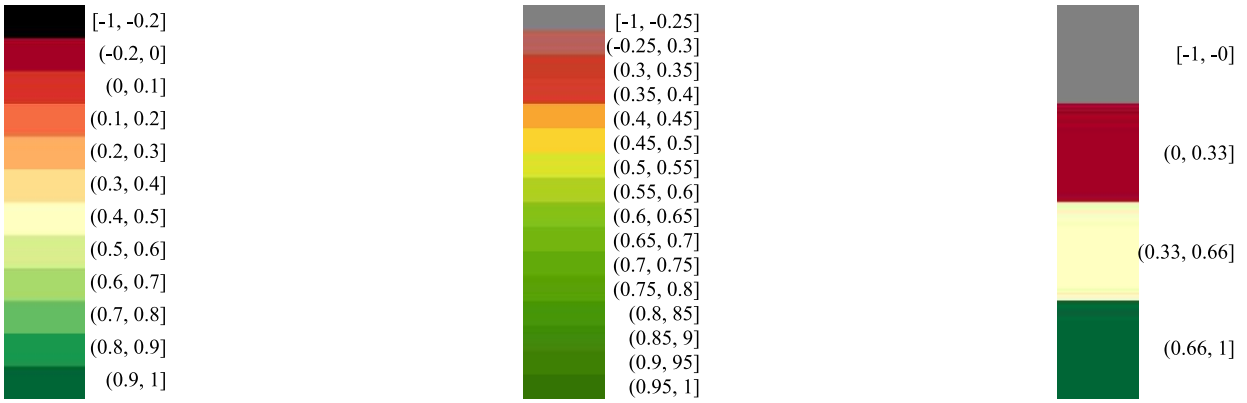

Figure 15: Sentinel palette (a), contrasting palette (b), four-colors palette (c).

NDVI color maps

The atmospheric correction of multi- and hyperspectral imagery is significant for accurate vegetation index computation. Research presented in [132] examines how different levels of atmospheric correction influence NDVI calculations.

Fig. 16 presents experimental results for both images across different product levels. The NDVI value distributions are illustrated in Fig. 17. For *PRISMA* data, all three computational strategies yield similar results, as evidenced by overlapping histograms. While Sentinel-2 data is available only with basic atmospheric correction, *PRISMA* imagery shows distinct patterns across processing levels. The L1 *PRISMA* histogram maintains a distribution similar to Sentinel-2, but L2B and L2C corrections shift the distribution toward higher NDVI values, suggesting healthier vegetation. This progression demonstrates how additional atmospheric corrections can significantly influence NDVI interpretation.

The increasing availability of hyperspectral data presents new opportunities for precision agriculture applications. Vegetation indices like NDVI provide farmers with crucial insights for crop management, enabling early detection of plant stress, optimization of irrigation and fertilization, and timely intervention in case of diseases. These vegetation maps serve as valuable decision-support tools for precision agriculture, helping farmers reduce costs while maximizing yield and sustainability.

The study [132] demonstrates that NDVI calculations vary significantly across different sensors and processing levels. While *PRISMA* L2C products tend toward optimistic NDVI estimates, Sentinel-2 L2A products show more conservative values. Although these three computational strategies for hyperspectral NDVI yielded consistent results within

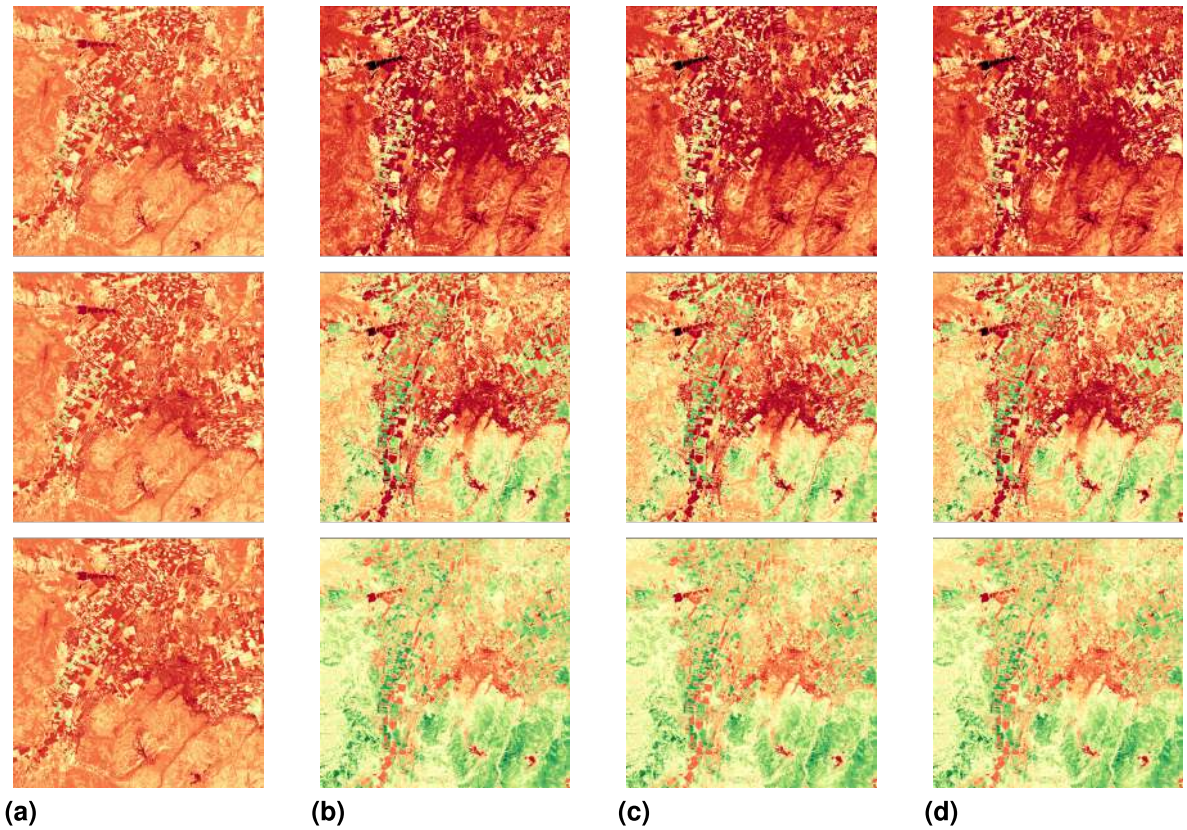

Figure 16: NDVI color maps for Sentinel-2 (column a) and *PRISMA* products using strategies (i) - column (b), (ii) - column (c) and (iii) - column (d) described in the NDVI section. In the top row are the results for L1, in the middle row from L2B and in the bottom row of L2C *PRISMA* products.

PRISMA data, the variation between sensors highlights the importance of ground validation for reliable agricultural decision-making. This validation is essential to ensure farmers receive accurate, actionable information for effective crop management strategies.

The paper [132] was presented at *IEEE WHISPERS* (Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing), a highly specialized conference dedicated to hyperspectral imaging, signal processing, and remote sensing applications. WHISPERS represents one of the premier international forums for presenting advances in hyperspectral data processing, analysis, and applications, bringing together experts from academia, industry, and space agencies to address challenges in hyperspectral remote sensing.

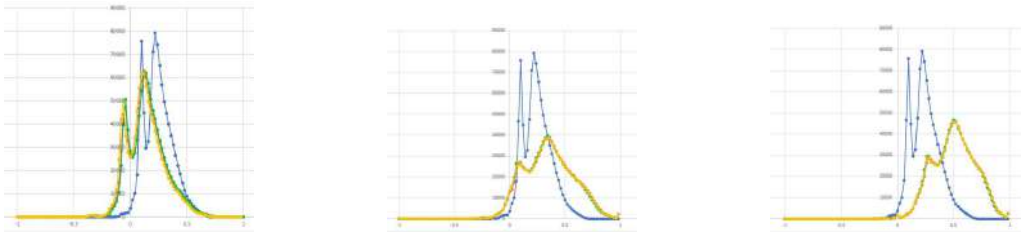

Figure 17: NDVI histograms for Sentinel-2 (blue color) and *PRISMA* L1 (a), L2B (b), L2C (c) using strategies (i) - red color, (ii) - green color, (iii) - yellow color.

I.3. Data quality and anomaly detection

Satellite imagery plays a crucial role in precision agriculture, requiring effective visualization and analysis methods. The study [130], co-authored by the author of the present thesis, examines how dataset characteristics influence ML-based visualization outcomes, finding that data quantity and variability significantly impact performance. The use of Macbeth color charts was explored to improve spectral analysis consistency, coupled with an efficient methodology for multispectral image interpretation. Using the NDVI, this approach enables automatic identification of potential crop development issues, offering practical benefits for agricultural health monitoring and early intervention strategies. This approach combines standardized color calibration techniques with automated analysis, providing practical tools for agricultural monitoring and decision-making.

Color image rendering and visual interpretation

As discussed earlier, multispectral and hyperspectral images cannot be directly displayed on standard digital devices. To enable human users to visually interpret these images and provide an initial layer of validation, it is necessary to apply a color mapping technique that converts the spectral data into a color RGB image that closely represents the natural scene captured by the sensor. An appropriate dataset is essential for applying a ML approach to image visualization. Such a dataset must include an RGB image for each spectral image in the set, enabling the mapping of spectral pixels to RGB triplets. Given the limited availability of datasets, the focus of the study [130] was on constructing a new, carefully designed dataset aimed at accurately representing the colors from the Macbeth color charts.

The Macbeth color chart serves as a standardized calibration target for mapping multispectral pixels to RGB colors. It comes in two variants: a basic 8-color chart representing RGB cube corners, and an enhanced 24-color version designed to match natural objects' spectral reflectances, including vegetation, skin tones, and flowers. The charts are represented in Fig. 18. The Macbeth charts' calibration to natural color spectra suggested their potential to improve visualization outcomes in the FCNN implementation [133].

In order to construct a dataset using the Macbeth charts, spectral images of the charts were acquired with a portable SPECIM IQ HS camera with a high spectral resolution (204 bands) in controlled laboratory conditions. The RGB visualization of the acquired images is presented in Fig. 18c and Fig. 18d. Two small datasets, one for each chart, were constructed from these HS images by manually selecting small areas inside each colored rectangle. The dataset from the 8-color chart contains 3,704 HS pixels and 8 labels (corresponding to the 8 colors), while the dataset from the 24-color chart contains 12.168 HS pixels and 24 labels (corresponding to the 24 colors). The labels were produced by specifying the exact RGB color coder corresponding to the selected rectangles. The labels for the 8-color chart are presented in Table 3, while the labels for the 28-color chart are illustrated in Table 4.

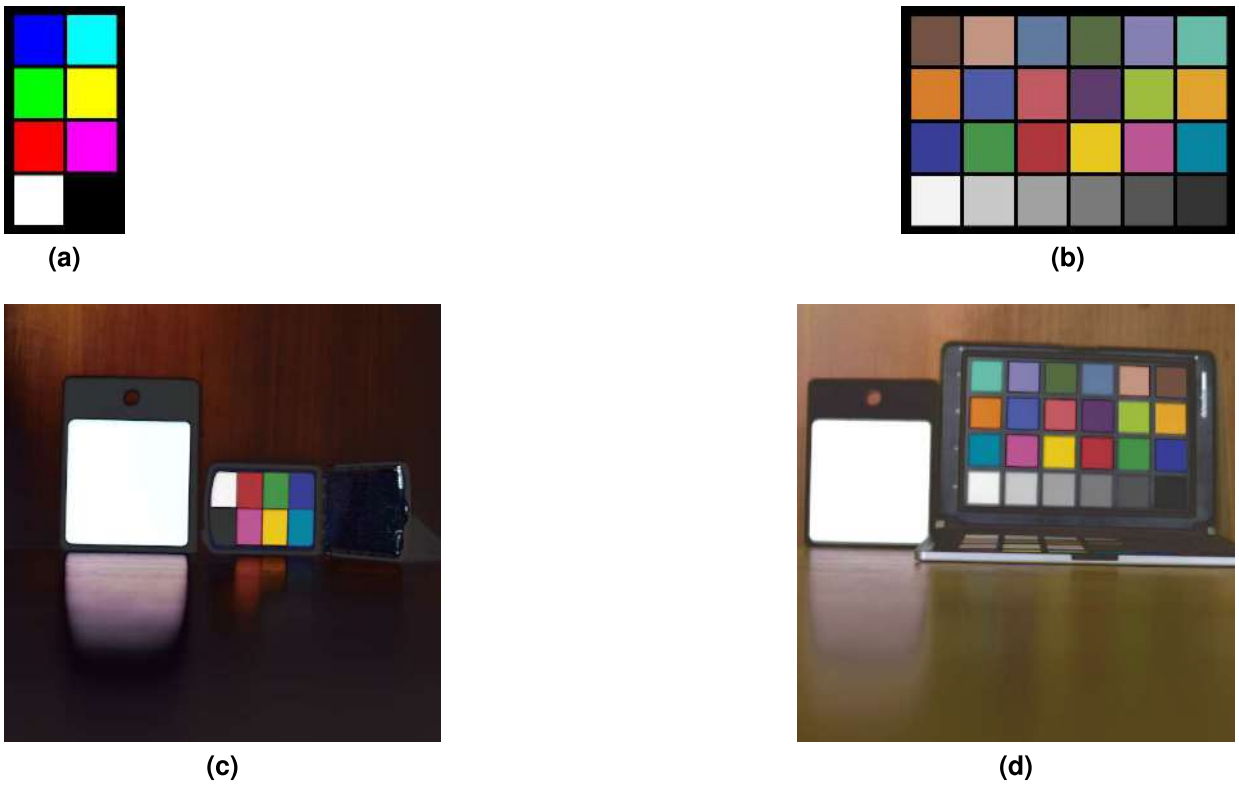

Figure 18: a) 8-color chart, b) 24-color Macbeth chart, c) RGB visualization of the 8-color-chart acquired hyperspectral image, d) RGB visualization of the 24-color-chart acquired hyperspectral image.

Table 3: Colors and corresponding labels of the 8-color chart.

Color				
RGB label	(0,0,255)	(0,255,255)	(0,255,0)	(255,255,0)
Color				
RGB label	(255,0,0)	(255,0,255)	(255,255,255)	(0,0,0)

The FCNN architecture from [133] was adapted to handle 204 spectral bands instead of the original 31 wavelengths. he network was trained separately on both 8-color and 24-color Macbeth charts, and its performance was evaluated on *PRISMA* satellite imagery and compared with classical visualization methods and previous FCNN results [133].

As can be seen from Fig. 19, the results of the network trained on the chart-datasets are not as qualitative, as expected. Although the spectral image samples and the corresponding labels were constructed with care, and the colors are very precise, the results are worse than in the case of the FCNN presented in [133], trained on the publicly available *CAVE* and *UGR* datasets. However, the small number of colors in the training set leads to a strong color quantization effect in the resulting image (which could be useful for a further compression and segmentation step). Additional results, including comprehensive image quality assessments (IQA), are presented in the original study [130].

Table 4: Colors and corresponding labels of the 24-color MabBeth chart

Color				
RGB label	(115,82,68)	(194,150,130)	(98,122,157)	(87,108,67)
Color				
RGB label	(133,128,177)	(103,189,170)	(214,126,44)	(80,91,166)
Color				
RGB label	(193,90,99)	(94,60,108)	(157,188,64)	(224,163,64)
Color				
RGB label	(56,61,150)	(70,148,73)	(175,54,60)	(231,199,31)
Color				
RGB label	(187,86,149)	(8,133,161)	(243,243,242)	(200,200,200)
Color				
RGB label	(160,160,160)	(122,122,122)	(85,85,85)	(52,52,52)

Analysis of anomalies in NDVI maps

Early detection of crop health issues remains crucial in precision agriculture. The study [130] leverages NDVI maps for anomaly detection in satellite imagery, identifying regions with unusual NDVI values that may indicate potential crop stress. Local neighborhood analysis of NDVI values allows the detection of subtle variations and potential issues even within apparently healthy agricultural zones.

In the study [130], co-authored by the author of the present thesis, anomalies are defined as pixels with lower NDVI values surrounded by healthy vegetation in 3×3 pixel vicinity. A distinction is made between moderate anomalies (Fig. 20a), where yellow pixels (NDVI: 0.33-0.66) are surrounded by healthy green pixels (NDVI: 0.66-1), and severe anomalies (Fig. 20b), where red pixels (NDVI: 0-0.33) are surrounded by healthy green pixels (NDVI: 0.66-1). This classification is based on the principle that agricultural crops typically develop uniformly, resulting in consistent NDVI values within local neighborhoods. Any deviation from this pattern may indicate potential crop health issues.

The hyperspectral image used in these experiments was obtained from the Italian Space Agency (ASI)'s *PRISMA* satellite on 24 March 2023 over a region of Braşov county. Images collected by the *PRISMA* satellite cover a range of 239 spectral bands, from 400 to 2500 nm, with 66 bands in the visible and near-infrared (VNIR) and 173 in the short wave infrared (SWIR) domains. The sampling interval for the spectral bands is smaller than 12 nm, thus providing a significant amount of information. The resolution of the *PRISMA* image is 1000×1000, with a spatial resolution of 30m. The *PRISMA* satellite provides 4 product levels of the raw image, with different types of correction. The product used in this study is the L2C level, comprising atmospheric correction.

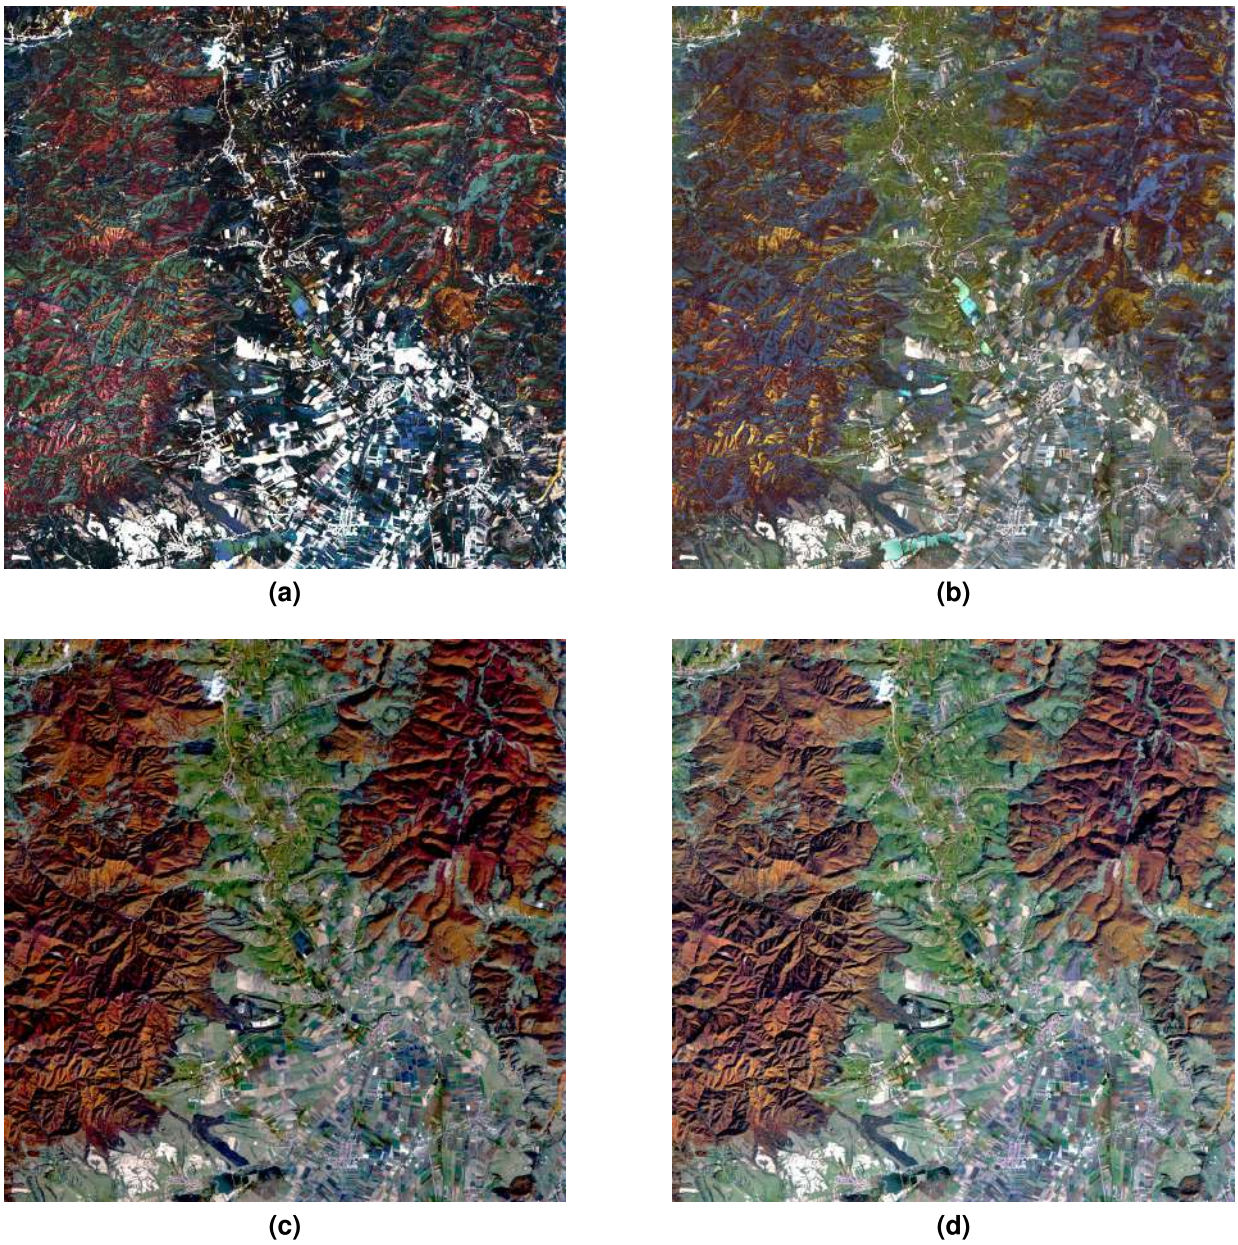

(a) **(b)**

(c) **(d)**

Figure 19: *PRISMA* HS image visualized with a FCNN trained: a) on 8-color chart; b) on 24-color chart; c) on *CAVE* dataset [133]; d) on *UGR* dataset [133].

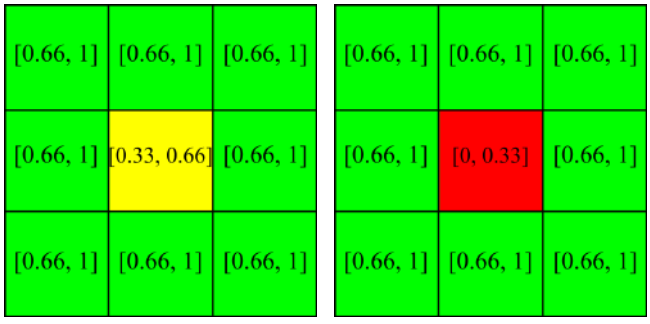

(a) **(b)**

Figure 20: Anomaly definition in NDVI maps.

Fig. 21a presents a visualization of the *PRISMA* image where this was performed by a FCNN trained on the *CAVE* dataset offering a significant color diversity in the visible domain [131].

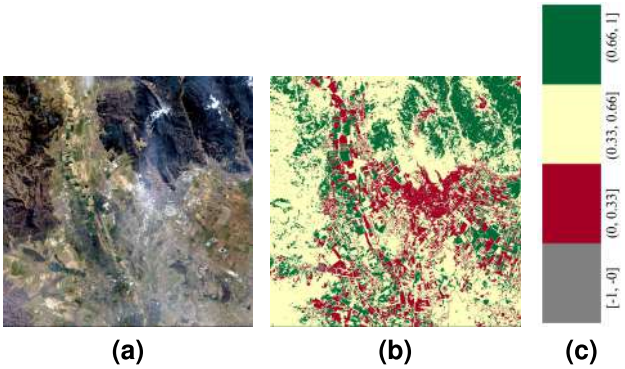

Figure 21: a) Original *PRISMA* image RGB visualization; b) its corresponding NDVI map; c) NDVI color palette.

The analysis identified 646 anomalies in the *PRISMA* image, with Fig. 22 presenting four detailed examples of agricultural parcels where detected anomalies are highlighted in magenta.

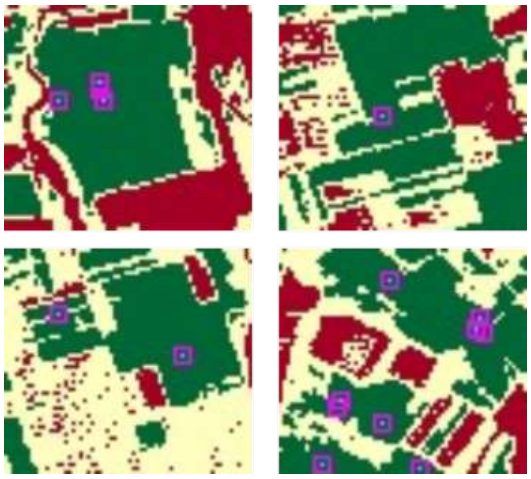

Figure 22: Zoom in image crops with detected anomalies (marked with magenta).

These experiments demonstrates that dataset quantity outweighs label precision in visualization performance, while also establishing an efficient NDVI-based methodology for crop anomaly detection. This approach enables rapid identification of agricultural issues, though further validation through temporal analysis and ground measurements is recommended. This research [130], covering both Macbeth color chart calibration tests for data quality assessment and anomaly detection methodologies, has been submitted and accepted to a Scopus-indexed journal, *Bulletin of the Transilvania University of Braşov. Series III: Mathematics and Computer Science*, and is currently in press.

II. Multisource data aggregation and dataset development for crop classification

In precision agriculture and remote sensing, harmonizing data from multiple satellite sources is a key step toward improving model performance and reducing noise variability across observation platforms. This section presents two complementary directions: the first explores methodological innovations for spectral image aggregation across different sensors, while the second describes the development of a new open-access dataset, DACIA5, designed specifically for agricultural crop classification using Sentinel-1 and Sentinel-2 imagery.

The proposed approaches emphasize the importance of integrating heterogeneous data in a unified framework, enabling downstream tasks such as vegetation analysis and crop-type prediction. DACIA5 serves not only as a research tool but also as a foundational asset for training and evaluating machine learning models in EO-based agricultural analytics.

By addressing the challenges of spectral compatibility, temporal alignment, and label consistency, the studies in this chapter provide essential infrastructure and methodological groundwork for robust AI-driven agricultural intelligence.

To support these approaches, this section offers a concise overview of selected research articles that apply and evaluate the methods discussed. The main findings and key contributions are briefly summarized here, with more detailed analysis provided in the subsequent subsections.

Selected studies and core insights:

- R.I. Luca, **A. Băicoianu**, I.C. Plajer, 2025 - *Spectral Image Data Aggregation for Multisource Data Augmentation* [101]

  Main findings:

  - The paper introduces a methodology for aggregating spectral image data, enabling compatibility and integration of spectral datasets from diverse sources.
  - It employs interpolation techniques to perform spectral data aggregation, significantly reducing the computational complexity and resource demands associated with traditional machine learning approaches (such as deep learning super-resolution, convolutional neural networks, and auto-encoders). As a result, the proposed method is both efficient and cost-effective, producing a unified dataset suitable for any machine learning application.
  - The approach uses six spectral datasets with differing spectral signatures, aggregating them based on the dataset with the lowest spectral resolution.
  - Four interpolation techniques (linear interpolation, quadratic interpolation, cubic spline interpolation, and piecewise cubic Hermite interpolating polynomial) were evaluated while highlighting the challenges in determining the most suitable one.

  - The study develops a comprehensive validation framework focusing on visual inspection (using 2D and 3D plots) and quantitative evaluation based on surface difference analysis, established metrics from existing literature, and custom-designed performance indicators. The framework also assesses the impact of the aggregation methodology on segmentation accuracy using two types of neural networks: a fully connected neural network and a U-Net architecture.
  - It proposes a custom evaluation metric for comparing interpolated and original spectral pixels. Based on the *Mean Squared Error*, the metric assumes that accurate interpolation preserves the original pixel shape, such that re-interpolating to the source wavelengths yields results nearly identical to the original pixel.
  - The approach also utilizes the NDVI to assess the effect of interpolation on vegetation index computation.

- **A. Băicoianu**, I.C. Plajer, M. Debu, M. Ștefan, M. Ivanovici, C. Florea,..., L. Dogar, 2025 - *DACIA5: a Sentinel-1 and Sentinel-2 dataset for agricultural crop identification applications* [20]

  Main findings:

  - Introduces a novel labeled dataset tailored for agricultural applications, comprising multispectral images from Sentinel-2 alongside synthetic aperture radar (SAR) data from Sentinel-1, collected over Brașov County spanning five years (2020–2025). It contains multispectral patches and crop-specific field parcel annotations, making it a strong foundation for training and validating models in real-world precision agriculture scenarios.
  - Provides highly accurate ground truth labels developed in collaboration with the National Institute of Research and Development for Potato and Sugar Beet, ensuring data reliability and field-validated labels.
  - Demonstrates practical crop identification case studies using both Sentinel-2 multispectral data and fused datasets combining Sentinel-2 and Sentinel-1 imagery, highlighting the benefits of data fusion for improved classification performance.
  - The section titled "Data set value" presents practical applications of the DACIA5 dataset through two case studies: one using a pretrained deep learning model (ResNet18), and the other employing a Random Forest classifier.

II.1. Aggregating spectral images from multiple sources

Problem statement and datasets

The increasing availability of MS and HS remote sensing data offers unique insight into material spectral characteristics, providing distinct spectral fingerprints beyond the

limited bands of RGB images. However, leveraging this data for machine learning tasks faces significant challenges, primarily due to inconsistencies in spectral signatures and spatial resolutions across different satellites and equipment. The scarcity of labeled data further restricts the practical use of these datasets, and traditional data augmentation methods [59] are often insufficient for very small datasets. This necessitates a preprocessing step to aggregate diverse spectral data into a uniform, consistent format, enabling broader application in machine learning models and facilitating generalization across different sensors. Currently, research in this area is limited and a standardized solution is lacking.

This section introduces a simple and cost-efficient methodology to integrate HS data from various sources into a single, larger dataset. This approach involves analyzing available datasets to identify similarities and differences, selecting a reference dataset, and performing interpolation operations over spectral bands to aggregate all images. To validate the utility, correctness, and accuracy of this methodology, two neural networks, a FCNN and a CNN, specialized in semantic segmentation, are trained and tested on the aggregated dataset. The networks are also tested on images from other sources, which were not used for the training.

The complete methodological approach and experimental validation of this technique is thoroughly documented in [101]. The proposed methodology begins with a comprehensive analysis of publicly available hyperspectral datasets, focusing on their spectral characteristics and suitability for integration. Selected datasets were then unified through interpolation techniques to create a consolidated spectral database.

For the experiments, several public HS datasets with diverse spectral signatures were utilized. These included *Pavia University*, *Kennedy Space Center (KSC)*, *Botswana*, and *Indian Pines* [66], all of which provide ground truth labels suitable for neural network classification tasks. To enable consistent testing across datasets, their original labels were combined into two categories: Vegetation and Non-Vegetation. Additionally, datasets without segmentation labels, such as *Columbia Imaging and Vision Laboratory (CAVE)* [191] and *Universidad de Granada (UGR*) [51], were incorporated for further analysis of interpolation methods.

The *Pavia University* image, captured by the ROSIS sensor, is a corrected dataset of 610×340 pixels with 103 spectral bands (430-860 nm) and a 4 nm spectral resolution. Originally featuring nine ground truth classes, these were processed into the two study-specific classes. Fig. 23 illustrates the original and processed labels.

The *KSC* image, captured by the AVIRIS sensor, contains 176 corrected spectral bands (400-2500 nm) with a 10 nm spectral resolution, over a 512×614 pixel area. Its original 13 ground truth classes were also consolidated into two. Fig. 24 displays these labels.

The *Botswana* image, collected by the Hyperion sensor, consists of 1476×256 pixels with 145 corrected spectral bands (400-2500 nm) and a 10 nm spectral resolution. Its 14

(a) (b) (c)

Figure 23: Pavia University: (a) Visualization using bands-selection; (b) Original ground truth; (c) Processed ground truth (*black* for unknown; *green* for vegetation; *red* for non-vegetation)

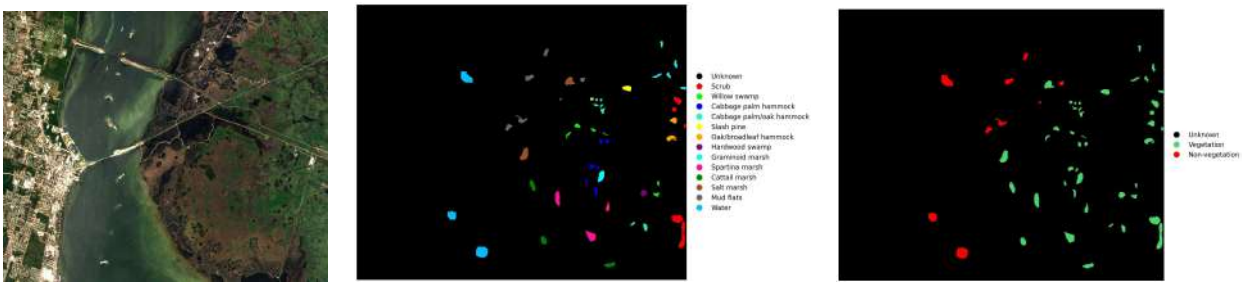

(a) (b) (c)

Figure 24: *KSC* image: (a) Visualization using band-selection; (b) Original ground truth; (c) Processed ground truth (*black* for unknown; *green* for vegetation; *red* for non-vegetation)

original labels were merged into two, as shown in Fig. 25.

The *Indian Pines* dataset, gathered by the AVIRIS sensor, is a 145×145 pixel image with 200 corrected spectral bands (400-2500 nm) and a 10 nm spectral resolution. Its initial 16 ground truth classes were also consolidated, as depicted in Fig. 26.

The *CAVE* dataset comprises 32 scenes (512×512 pixels, 31 bands, 400-700 nm, 10 nm resolution), while the *UGR* dataset contains 14 outdoor scenes (1000×900 pixels, 61 bands, 400-1000 nm, 10 nm resolution). Neither dataset provides segmentation labels. These datasets were also used in the study [133], co-authored by the author of the present thesis.

The *Hyperview* dataset [119], with its high spectral resolution, holds potential as a reference for interpolation but was not the primary focus of this study. The *CloudSEN12* dataset [9], designed for cloud detection, was excluded as its focus does not align with hyperspectral image classification for vegetation. From the *Onera Satellite Change Detection Dataset (OSCD)* [43], primarily for change detection, two specific images (*Norcia* and *Chongqing*) were used to demonstrate the extensibility of the methodology. These

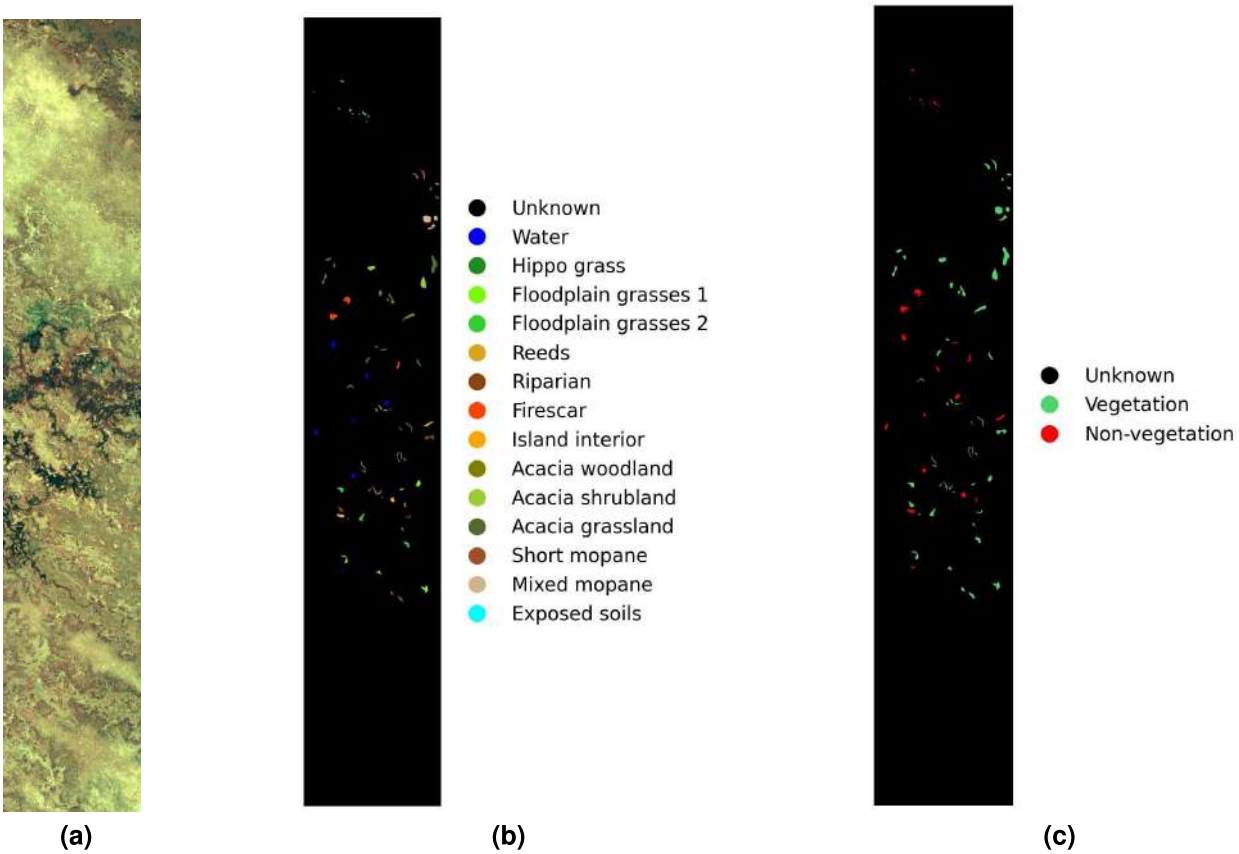

Figure 25: *Botswana* image: (a) Visualization using band-selection; (b) Original ground truth; (c) Processed ground truth (*black* for unknown; *green* for vegetation; *red* for non-vegetation)

Figure 26: *Indian Pines* image: (a) Visualization using band-selection; (b) Original ground truth; (c) Processed ground truth (*black* for unknown; *green* for vegetation; *red* for non-vegetation)

selected images are depicted in Fig. 27.

Spectral interpolation methods

Spectral data aggregation from multiple HS sources often yields inconsistent formats, posing a challenge for traditional neural networks. The proposed approach addresses this by employing a preprocessing step that uses interpolation to transform data into a predetermined format with a consistent number of channels and wavelengths, thereby potentially increasing the information content per pixel.

(a) **(b)**

Figure 27: *OSCD* Images: (a) Norcia; (b) Chongqing

For this aggregation, the *Pavia University* scene, with its 4 nm spectral resolution, was chosen as the reference for interpolation. To avoid unrealistic results, particularly for datasets like *CAVE*, wavelengths were limited to a maximum of 690 nm. Values for each *Pavia University* channel are interpolated from adjacent channels of other processed images. Four classical interpolation techniques, well-known in various domains, were applied and compared for their effectiveness in aggregating spectral images.

The following interpolation techniques were evaluated and documented in this section, each offering distinct advantages in harmonizing spectral data across different hyperspectral sources: linear interpolation, quadratic interpolation, cubic spline interpolation, Piecewise Cubic Hermite Interpolating Polynomial (PCHIP). For detailed mathematical definitions and implementation specifics of these methods, see [101].

Methodology, experimental results, and comparative quality assessment

Two primary strategies were employed to validate the proposed method and assess the interpolation results. The first involved visual inspection combined with selected quantitative metrics, while the second focused on evaluating how the methodology influenced segmentation accuracy when using neural networks.

Evaluation of interpolation quality through various metrics

The initial assessment of the spectral pixel interpolation results was conducted through visual analysis, utilizing 2D and 3D comparative plots, as detailed below.

The interpolation results for each image were compared both to the original reference and to one another. This was achieved through two distinct types of graphical representations: 2D plots of individual pixels and 3D surface plots of the entire image. These visualizations help identify which interpolation method more accurately preserves the shape

of the original spectral distribution. While these visualizations offer valuable preliminary indicators of method suitability, they cannot solely determine interpolation effectiveness, as the primary objective extends beyond preservation of original information to include information content enhancement. The 2D analysis focused on random pixel selections, comparing original spectral signatures against their interpolated counterparts. Representative results are presented in Fig. 28 and Fig. 29, showing spectral signatures from *KSC* and *Botswana* datasets respectively.

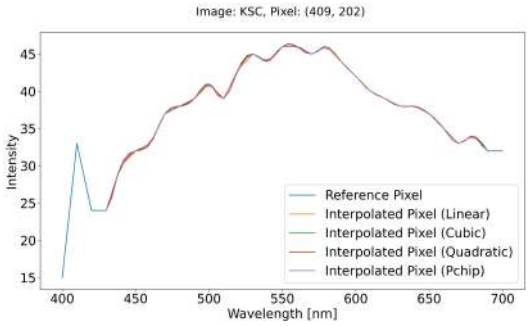

Figure 28: Reference vs interpolated Pixel for *KSC*.

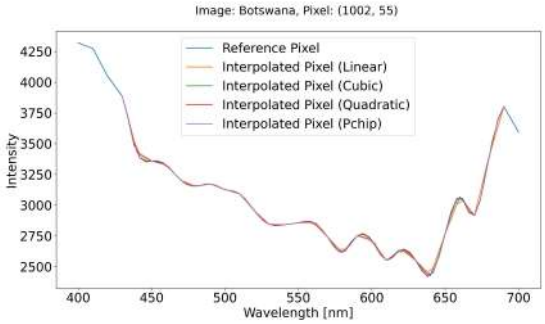

Figure 29: Reference vs interpolated Pixel for *Botswana*.

The 3D plots offer a comprehensive visualization of the pixel surface distribution across the entire image, including multiple spectral bands. Some plots can be seen in Fig. 30 (one image of *CAVE* dataset), and Fig. 31 (*OSCD Chongqing*)

The graphical analysis demonstrates that all implemented interpolation methods adequately maintain the general spectral signature patterns, despite some variations in intensity values. The practical significance of these variations is further assessed through comprehensive neural network validation tests.

While numerous metrics exist for assessing HS image quality, such as Spectral Discrepancy, MSE, and SAM, their utility is limited when comparing images with differing numbers of spectral bands [64, 40]. In this case, where the number of spectral bands is increased through interpolation (from 10 nm to 4 nm resolution), a different validation approach was adopted.

To quantitatively evaluate interpolation quality, the surface under the 2D plots was computed using the Trapezoidal Rule [58]. This method offers a more precise comparison

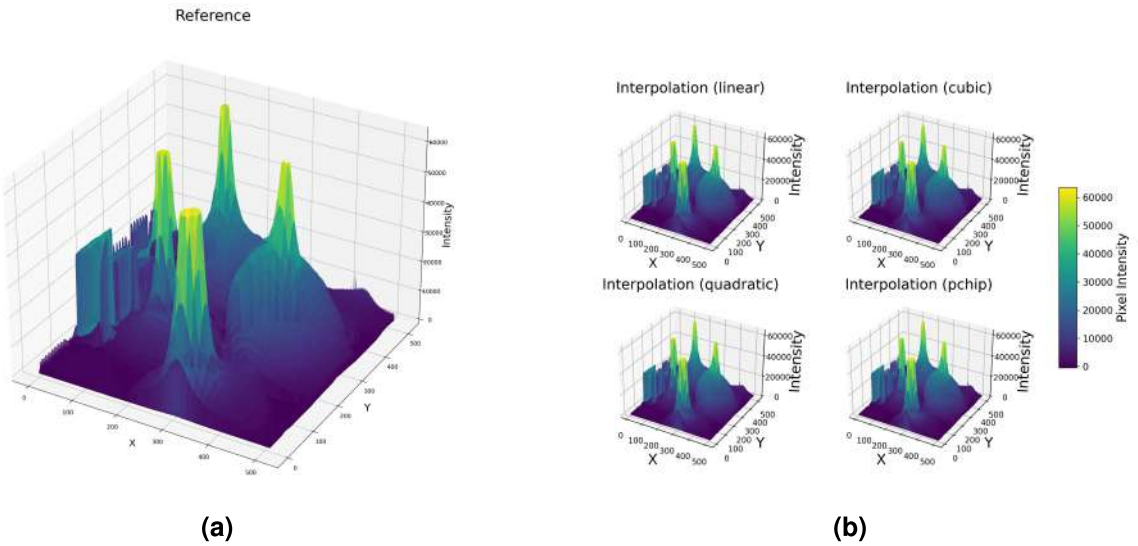

(a) **(b)**

Figure 30: Pixel surface for *CAVE* Balloons.

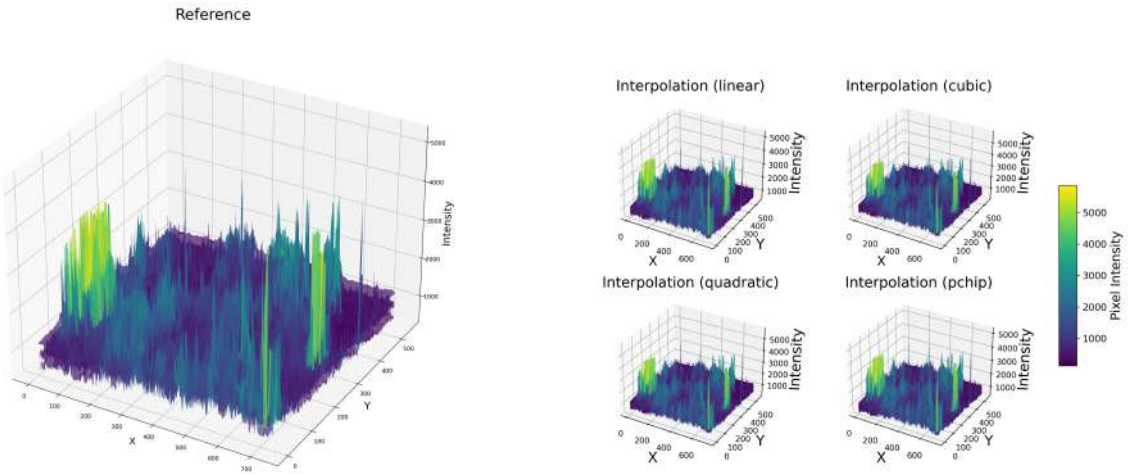

(a) **(b)**

Figure 31: Pixel surface for *OSCD* Chongqing.

than visual inspection, enabling the identification of interpolation techniques unsuitable for specific datasets based on extreme average differences from the reference. A summary of these results is presented in Table 5.

In addition, to consistently evaluate and compare different interpolation methods, a suitable metric was required, as existing plot-based measures and standard MSE were not directly applicable due to changes in spectral bands and the absence of ground truth. Therefore, a *Custom Mean Squared Error (CMSE)* measure was elaborated to assess how well an interpolated pixel preserves the shape of the original. For each MS pixel, it calculates the MSE between this pixel and the result of its interpolation onto the desired wavelengths and then back onto the original ones. The formula is expressed by Equation 3.

$$CMSE(p) = \frac{(p - I_b(I_f(p))^2}{nw} \tag{3}$$

Table 5: Surface average differences

Dataset	Interpolation	Surface Average Difference
CAVE Balloons	Linear	3.60
	Cubic	3.56
	Quadratic	3.60
	PCHIP	3.56
UGR	Linear	3.90
	Cubic	3.89
	Quadratic	3.88
	PCHIP	3.90
Indian Pines	Linear	1.89
	Cubic	1.88
	Quadratic	1.95
	PCHIP	2.13
KSC	Linear	6.01
	Cubic	7.32
	Quadratic	6.53
	PCHIP	5.54
Botswana	Linear	5.90
	Cubic	5.80
	Quadratic	5.79
	PCHIP	5.87
OSCD Norcia	Linear	74.54
	Cubic	79.26
	Quadratic	77.14
	PCHIP	76.70
OSCD Chongqing	Linear	137.46
	Cubic	134.77
	Quadratic	135.39
	PCHIP	138.79

in which $I_f(p)$ represents the *forward interpolation*, i.e. the interpolation of pixel p on the new wavelengths and $I_b(p)$ represents the *backward* interpolation of p, i.e. the interpolation onto the source wavelength, and nw the number of wavelengths.

The results, summarized in Table 6, indicate that linear interpolation yields a very large CMSE, while quadratic and cubic interpolations result in relatively small values. This outcome is expected, given that linear interpolation typically provides the roughest approximation for a signal.

Moreover, the influence of interpolation methods on vegetation indices calculation was thoroughly investigated [101]. The analysis revealed that all tested interpolation techniques demonstrated comparable suitability for spectral data processing. This significant finding suggests that linear interpolation, being the most straightforward and computationally efficient method, can be effectively employed for vegetation analysis in precision agriculture applications.

Evaluation of interpolation quality through the accuracy of the semantic segmentation

This section evaluates the impact of the proposed spectral data aggregation method on semantic segmentation accuracy for HS images. Two neural networks, a FCNN and a U-Net, were implemented for pixel-wise classification. Both models were trained to distinguish between "vegetation" and "non-vegetation", a classification chosen to align with agricultural applications and to unify disparate class definitions across datasets.

The implemented neural network architecture consists of an input layer with 66 neu-

Table 6: CMSE: Reference - Interpolated Interpolation

Dataset	Interpolation	MSE
CAVE Balloons	Linear	3675.01
	Cubic	**7.75**
	Quadratic	**5.27**
	PCHIP	341.31
UGR	Linear	7800.99
	Cubic	**23.12**
	Quadratic	**15.08**
	PCHIP	897.71
Indian Pines	Linear	554.50
	Cubic	**1.65**
	Quadratic	**1.09**
	PCHIP	78.59
KSC	Linear	96.17
	Cubic	**0.58**
	Quadratic	**0.37**
	PCHIP	8.76
Botswana	Linear	311.14
	Cubic	**0.84**
	Quadratic	**0.57**
	PCHIP	35.75
OSCD Norcia	Linear	24.48
	Cubic	**0.000035**
	Quadratic	$\approx$**0**
	PCHIP	0.04
OSCD Chongqing	Linear	13.83
	Cubic	**0.000012**
	Quadratic	$\approx$**0**
	PCHIP	0.02

rons (matching the preprocessed wavelength count), four fully connected hidden layers (128, 256, 512, and 256 neurons), and a two-neuron output layer for vegetation/non-vegetation classification. ReLU activation was used for the hidden layers, with Softmax applied to the output layer. The network training utilized the *Pavia University* dataset along with various other datasets previously described. The training process employed Cross-Entropy Loss and Adam optimizer (learning rate: 0.0001) over 150 epochs, with data randomly split into training (80%) and testing (20%) sets, processed in 2048-pixel batches.

The U-Net architecture, known for its U-shape, features a contracting (encoder) and an expanding (decoder) path, each comprising five 3D convolutional layers, followed by Batch Normalization and ReLU activation. The final decoder layer is a two-filter convolutional layer with Softmax. High-resolution features from the encoder are concatenated with the decoder's output for enhanced localization. Input patches of 10×10 pixels with 66 channels (shape: (10, 10, 66, 1)) were used. Filter counts for the contracting path were 64, 64, 128, 256, 256, and for the expanding path, 64, 128, 256, 512, 256. .

Training was performed on 10×10 patches from *Pavia University* and other datasets, with data split into 80% training, 10% validation, and 10% testing. Depending on the selected datasets, the number of epochs varies because the training is stopped when a monitored metric, the loss in this case, has stopped improving. Adam was employed as the optimizer with a learning rate of 0.001. Data augmentation [45], including horizontal/vertical flips and random rotations, was applied to training patches.

Four datasets (*Pavia University*, *KSC*, *Botswana*, *Indian Pines*) with merged vegeta-

tion vs non-vegetation labels were used for training and testing. The results for the FCNN and U-Net are in Table 7 and Table 8, respectively. Notably, Table 7 reveals that training with non-shuffled pixels led to overfitting, characterized by high training accuracy (exceeding 93%) but significantly poorer performance on new test datasets. This demonstrates that shuffling pixels during training is crucial for mitigating overfitting in the context of the employed neural network architecture and datasets.

Table 7: FCNN Training and Testing results

Training Dataset	Training Accuracy	Testing Dataset	Testing Accuracy
Pavia University	71.5/**94.15%**	KSC (Linear)	72.56/**57.59%**
		KSC (Cubic)	72.37/**58.34%**
		KSC (Quadratic)	72.37/**58.49%**
		KSC (PCHIP)	72.56/**58.28%**
		Botswana (Linear)	53.17/**25.46%**
		Botswana (Cubic)	53.48/**25.46%**
		Botswana (Quadratic)	53.17/**25.46%**
		Botswana (PCHIP)	53.17/**25.46%**
		Indian Pines (Linear)	90.75/**7.27%**
		Indian Pines (Cubic)	90.62/**7.62%**
		Indian Pines (Quadratic)	90.61/**7.49%**
		Indian Pines (PCHIP)	90.64/**7.40%**
Pavia University and Botswana (Linear)	88/**93.38%**	KSC (Linear)	72.56/**60.74%**
		KSC (Cubic)	72.40/**61.14%**
		KSC (Quadratic)	72.54/**61.12%**
		KSC (PCHIP)	72.56/**60.81%**
		Botswana (Linear)	90.02/**97.04%**
		Botswana (Cubic)	89.93/**96.86%**
		Botswana (Quadratic)	89.93/**96.95%**
		Botswana (PCHIP)	89.99/**96.98%**
		Indian Pines (Linear)	94.98/**7.94%**
		Indian Pines (Cubic)	94.84/**8.10%**
		Indian Pines (Quadratic)	94.89/**8.13%**
		Indian Pines (PCHIP)	94.97/**8.10%**
Pavia University and KSC (Cubic)	89.27/**93.28%**	KSC (Linear)	97.08/**96.49%**
		KSC (Cubic)	97.12/**96.49%**
		KSC (Quadratic)	97.12/**96.49%**
		KSC (PCHIP)	97.08/**96.49%**
		Botswana (Linear)	77.46/**25.46%**
		Botswana (Cubic)	77.77/**25.46%**
		Botswana (Quadratic)	77.83/**25.46%**
		Botswana (PCHIP)	78.14/**25.46%**
		Indian Pines (Linear)	92.97/**4.68%**
		Indian Pines (Cubic)	92.46/**4.69%**
		Indian Pines (Quadratic)	92.54/**4.69%**
		Indian Pines (PCHIP)	92.87/**4.69%**

Vegetation classification was conducted using the FCNN trained on *Pavia University* and *Botswana* linearly interpolated data, despite the *OSCD* dataset lacking classification labels, which precludes quantitative evaluation. The classification results, where vegetation is represented in green and non-vegetation in black, are shown in Fig. 32 for *Norcia* (cubic interpolated) and *Chongqing* (quadratic interpolated) images. For a comprehensive analysis of the results, including detailed precision metrics and F1 Score values across different training configurations and testing datasets, see [101].

The validation results, obtained through multiple testing approaches and neural network implementations, confirm the effectiveness of the proposed spectral data aggregation methodology. Despite using basic network architectures and working with inherent labeling constraints, the method demonstrated robust performance across diverse datasets. This versatility, combined with the method's adaptability, indicates significant

Table 8: U-Net Training and Testing results

Training Dataset	Training Accuracy	Testing Dataset	Testing Accuracy
Pavia University (82 epochs)	89%	KSC (Linear)	72.08
		KSC (Cubic)	
		KSC (Quadratic)	
		KSC (PCHIP)	
		Botswana (Linear)	91.70
		Botswana (Cubic)	88.98
		Botswana (Quadratic)	90.52
		Botswana (PCHIP)	90.36
		Indian Pines (Linear)	69.41
		Indian Pines (Cubic)	71.19
		Indian Pines (Quadratic)	71.44
		Indian Pines (PCHIP)	71.13
Pavia University and Indian Pines (Linear) (95 epochs)	89%	KSC (Linear)	72.08
		KSC (Cubic)	
		KSC (Quadratic)	
		KSC (PCHIP)	
		Botswana (Linear)	86.11
		Botswana (Cubic)	86.10
		Botswana (Quadratic)	86.11
		Botswana (PCHIP)	86.11
		Indian Pines (Linear)	94.96
		Indian Pines (Cubic)	94.95
		Indian Pines (Quadratic)	94.96
		Indian Pines (PCHIP)	94.96
Pavia University and KSC (Cubic) (89 epochs)	86%	KSC (Linear)	72.08
		KSC (Cubic)	
		KSC (Quadratic)	
		KSC (PCHIP)	
		Botswana (Linear)	92.43
		Botswana (Cubic)	92.44
		Botswana (Quadratic)	92.43
		Botswana (PCHIP)	92.43
		Indian Pines (Linear)	80.04
		Indian Pines (Cubic)	78.79
		Indian Pines (Quadratic)	78.70
		Indian Pines (PCHIP)	78.83

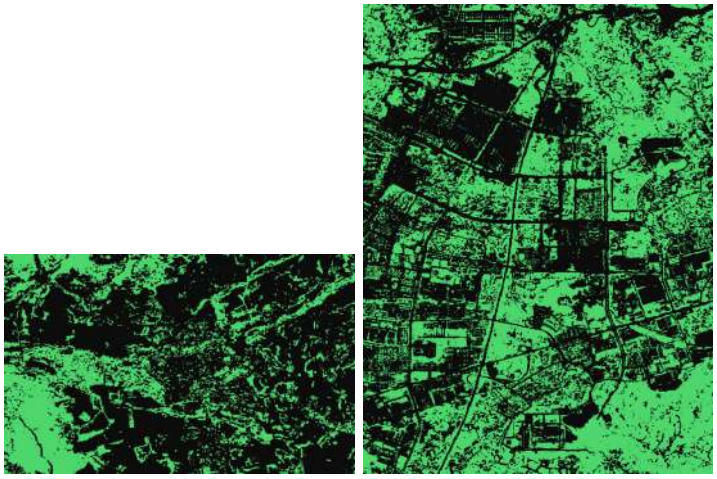

(a) **(b)**

Figure 32: *OSCD* classification results: (a) Norcia (Cubic); (b) Chongqing (Quadratic).

potential for various applications, particularly when coupled with advanced architectures or enhanced labeling techniques. The complete results presented in this section (II.1), initially developed as part of a master's thesis research, have evolved into a comprehensive study published in the *European Journal of Remote Sensing* [101], a Q2-ranked journal specialized in remote sensing applications.

II.2. Exploring the potential of the DACIA5 dataset for agricultural crop identification

Data quality challenges in ML for precision agriculture

Quality data form the foundation of any successful machine learning application in agriculture (and not only). Unlike other domains, agricultural data is often heterogeneous, collected under variable environmental conditions, and influenced by factors such as seasonality, sensor precision, and human annotation inconsistencies. Poor data quality - including issues like missing values, noise, low resolution, mislabeling, and sampling bias - can significantly impair model accuracy, generalization, and real-world applicability.

Before developing models, it is essential to ensure that the data meets key criteria: it should be labeled with ground truth, balanced to prevent bias, diverse in representing various agricultural conditions, high-resolution to capture subtle variations, current with modern farming practices, sufficiently large for generalization, standardized for reproducibility, and well-structured for efficient processing.

However, labeled datasets meeting these criteria are scarce in agricultural applications. This scarcity motivated the collaboration with the the National Institute of Research and Development for Potato and Sugar Beet from Brașov to develop a comprehensive dataset combining Sentinel-2 multispectral imagery and Sentinel-1 radar data, namely DACIA5. Covering agricultural areas north of Brașov from 2020 to 2024, this dataset provides field-validated labels and multispectral patches, creating a robust foundation for precision agriculture applications. The proposed dataset comprises 172 Sentinel-2 multispectral images (800×450 pixels) and 159 Sentinel-1 radar images. From these datasets, 6,454 Sentinel-2 and 5,995 Sentinel-1 patches (32×32 pixels) were extracted, totaling over 6 million labeled pixels. The data covers agricultural parcels managed by the National Institute of Research and Development for Potato and Sugar Beet, ensuring accurate field-validated labeling with detailed crop type information and phenological data.

DACIA5 multisource dataset development for crop classification

Before delving into the technical details and construction of the DACIA5 dataset, it is important to highlight that all the information discussed in this section, as well as the dataset contents themselves, are publicly accessible and can be explored. The DACIA5 dataset has been made available to the research community to promote transparency, reproducibility, and further advancements in agricultural remote sensing. Interested users can verify and visualize the dataset directly through the official repository https://zenodo.org/records/14915950, where the complete dataset is hosted along with accompanying documentation and tools.

Remote sensing, particularly satellite imagery, represents a major source of agricultural big data. The Copernicus Earth Observation Program, through its Sentinel constellation, provides extensive, freely available environmental data worldwide [8].

The Sentinel-2 satellites, operational since 2015-2017, deliver 13 spectral bands at resolutions between 10 m and 60 m [62]. Operating in Sun-synchronous orbit at 786 km altitude with a 290 km swath width, they ensure consistent imaging conditions. The DACIA5 dataset comprises Sentinel-2 Level-2A products from 2020-2024, covering the Research and Development Institute of *Transilvania University* of Brașov (45.669410°N, 25.549550°E).

All bands were resampled to 10×10m resolution using SNAP's Nearest Neighbour algorithm [161]. Image selection employed a 30% cloud coverage threshold through the Copernicus Browser platform, followed by manual verification of area clarity. In Table 9 is presented an overview of the spectral bands specific to Sentinel-2 images as described in [136]. It should be noted that Band 10 was excluded from the level 2A products during the process of atmospheric correction.

Table 9: Sentinel-2 bands with the central wavelengths (in nanometers) and their description

Band	Spatial Resolution (m)	Central Wavelength (nm)	Description
B1	60	443	Aerosol
B2	10	490	Blue
B3	10	560	Green
B4	10	665	Red
B5	20	705	Vegetation Red Edge 1
B6	20	740	Vegetation Red Edge 2
B7	20	783	Vegetation Red Edge 3
B8	10	842	NIR
B8A	20	865	Narrow NIR
B9	60	945	Water Vapor
B10	60	1375	Cirrus Clouds
B11	20	1610	SWIR1
B12	20	2190	SWIR 2

To extract and save parcel coordinates, an image was selected in which parcel boundaries were clearly visible. Delineation and coordinate extraction were performed using the *roipoly* function in MATLAB, applied to summer-acquired imagery, when parcels are most distinctly observable. The process involved manually selecting the parcel corners in the RGB image, with the corresponding coordinates saved in *.txt* files. Once all parcels identified on the National Institute of Research and Development for Potato and Sugar Beet's maps were digitized, binary masks were generated for each parcel across different years. These masks were derived from the *.txt* files, which contained the spatial definitions necessary for mask creation. The number of such files is summarized in Table 10.

Table 10: Overview of acquired Sentinel-1 and Sentinel-2 images from DACIA5 dataset

Year	No images Sentinel-1	No images Sentinel-2	No txt files	No crops	No patches Sentinel-1	No patches Sentinel-2
2020	40	40	51	10	1520	1520
2021	32	32	53	12	1380	1380
2022	25	30	47	12	914	1121
2023	28	30	47	12	1289	1410
2024	34	40	48	11	892	1023

The Sentinel-1 mission, part of ESA's Copernicus program, provides all-weather, day-and-night SAR images operating at 5.405 GHz (C-band). SAR uses satellite motion

to synthetically increase aperture for high-resolution imaging. Backscatter depends on soil and vegetation structure and moisture.

In DACIA5 dataset, Sentinel-1 images temporally close to Sentinel-2 images were used, selecting SAR acquisitions within a three-day window around each Sentinel-2 date. Data were retrieved from Google Earth Engine's `COPERNICUS/S1_GRD` collection, acquired in Interferometric Wide (IW) swath mode with VV and VH polarizations at 10 m resolution.

The images were preprocessed using the Sentinel-1 Toolbox, including detection, multilooking, ground projection (WGS84), noise removal, calibration, terrain correction (with SRTM or ASTER DEM), reprojection to UTM, and conversion to decibels. A Lee filter was applied to reduce speckle noise. Fig. 33 illustrates the processing workflow. More details about all the processing steps can be found in [20].

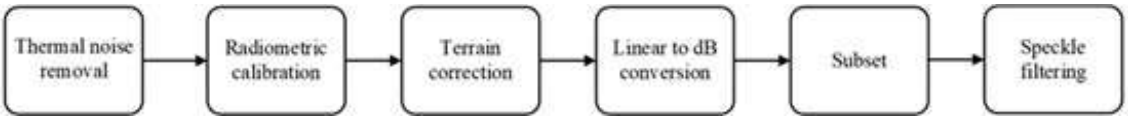

Figure 33: Processing steps for the Sentinel-1 data using the Sentinel-1 Toolbox).

A MATLAB script was used to generate masks from coordinates stored in *.txt* files, organized in separate directories by year. While parcel coordinates remain mostly unchanged, crops vary annually. Masks were refined using data from the National Institute of Research and Development for Potato and Sugar Beet. Each year, parcels were reviewed against updated maps to correct any error in positioning or crop labeling.

The mask files follow a naming convention: `a_b`, where `a` is the parcel number and `b` is the unique crop code assigned by the Payments and Intervention Agency for Agriculture (APIA) in Romania (see Table 2 from [20]). For example, the file named `1_151.mat` contains multispectral data from parcel 1 with crop code 151, corresponding to green peas. Each year, masks were generated for each parcel with labels matching that year's crops. The masks' visual representation allows easy identification of crops cultivated by the National Institute of Research and Development for Potato and Sugar Beet, using a color palette from the United States Department of Agriculture. These colored masks reveal crop distribution, diversity, and rotation over the years. Fig. 34 shows the resulting masks. The National Institute of Research and Development for Potato and Sugar Beet has dedicated parcels for research studies, and these parcels remain the same annually, but crops change based on study needs. Also, note that not all 17 crops appear every year.

In Table 10 are specified the number of Sentinel-1 and Sentinel-2 images acquired in each year from 2020 to 2024, the number of text files containing the coordinates of the parcels, the number of crops in each year, as well as the number of patches generated from the Sentinel-1 and respectively the Sentinel-2 images. The number of text files differs in the different years, as some of the parcels were divided in some years into subparcels.

To provide more details about the data records, the DACIA5 database contains all available Sentinel-2 MS images and corresponding Sentinel-1 SAR data spanning five

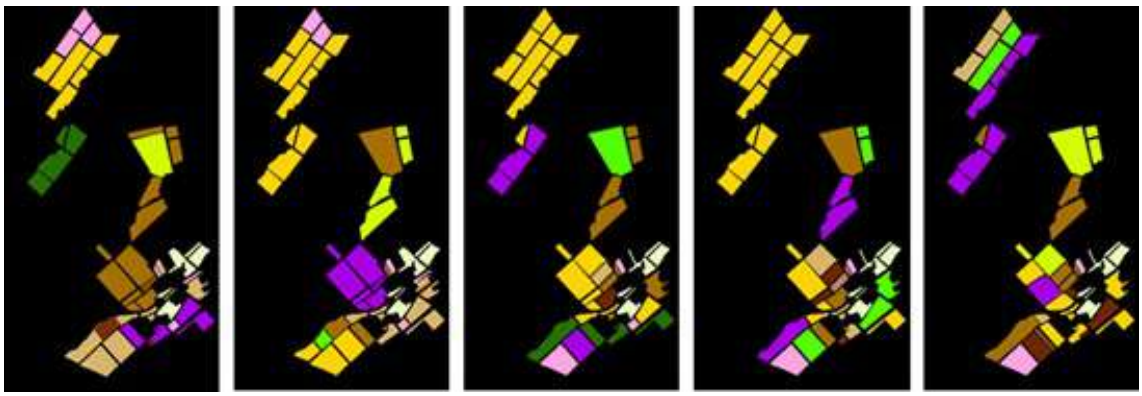

Figure 34: Processing steps for the Sentinel-1 data using the Sentinel-1 Toolbox).

years (2020–2024) over the area of interest. All images follow a naming convention and are saved in GeoTIFF format, compatible with MATLAB and Python. The naming scheme is detailed in Fig. 35.

Figure 35: Image naming convention. YYYY - is the four-digit year, MM - is the two-digit month, and DD - is the two-digit day.

Sentinel-2 images have dimensions of $800 \times 450 \times 12$ (height × width × spectral bands) with 10×10 m spatial resolution. Sentinel-1 images measure $800 \times 450 \times 2$ with the same spatial resolution. Sentinel-1 data are stored in the folder `Images_Sentinel1_2bands_GeoTIFF`, while Sentinel-2 data reside in `Images_Sentinel2_12bands_GeoTIFF`. Inside each folder, images are organized by year into subfolders named `Sentinel1_yyyy` and `Sentinel2_yyyy`, where `yyyy` is the acquisition year. The directory `Images_Sentinel2_GeoTIFF` contains Sentinel-2 images with full Copernicus metadata. Table 3 summarizes the number of images per year.

Besides Sentinel-1 and Sentinel-2 data, the database includes ground truth agricultural crop masks in PNG format and labeled masks in PNG and MAT formats, stored under `Masks_and_legend`. This directory also contains a PDF legend and yearly subfolders `Masks_yyyy` with masks for each year.

Using the Sentinel-1 and Sentinel-2 imagery, SAR and multispectral patches of size $32 \times 32 \times 2$ and $32 \times 32 \times 12$ were generated, respectively, stored systematically in the `32x32_patches` folder.

Patch selection was performed by sliding a 32×32 pixel window over the images from top-left to bottom-right with a step of 16 pixels. Simultaneously, the corresponding yearly label masks were checked to count pixels of each crop within the patches. A patch was considered valid if at least 75% (768 of 1024) pixels belonged to a single crop. Valid patches were saved following a specific naming convention.

Fig. 36a shows the traversal positions and resulting patches. Fig. 36b and Fig. 36c present overlays of patches on Sentinel-1 and Sentinel-2 images, respectively, demonstrating that patches correspond spatially across both sensors. Each patch is named according to a structured convention, as exemplified in Fig. 37.

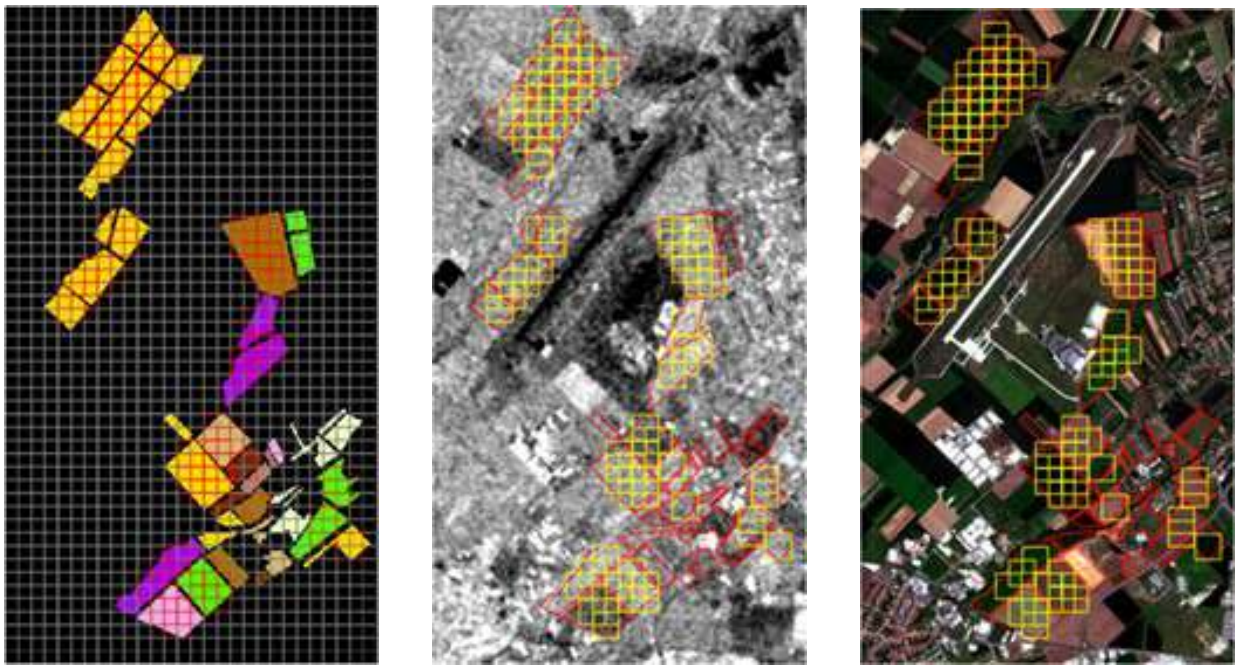

(a) (b) (c)

Figure 36: Selection grid for the patches.

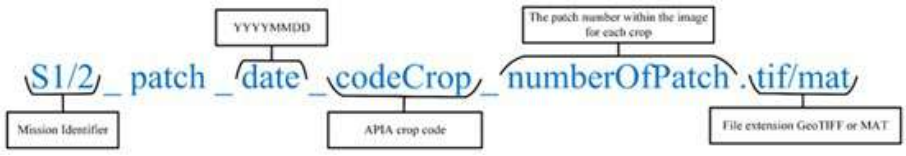

Figure 37: The naming convention for patches.

The `32x32_patches` directory contains two subfolders: `32x32_SAR+MSI_patches` with radar and multispectral data for two tasks - ACI (agricultural crop identification with temporal generalization) and ACEI (early crop identification) - and `32x32_RGB_patches` for visualization using RGB masks as ground truth. Data is split into training and test sets. Patches are georeferenced and saved in both GEOTIFF and MAT formats, in folders named `sentinel1_patches_mat`, `sentinel1_patches_tiff`, `sentinel2_patches_mat`, and `sentinel2_patches_tiff`.

The `32x32_RGB_patches` folder contains patches from RGB masks, where each patch must contain at least 75% of a single crop to be valid. Yearly subfolders named `patches_yyyy` organize the masks by year.

The `RoI_and_labels` directory stores geographic information and crop labels, divided into three subfolders: `RoI` (region of interest polygon), `WGS84`, and `UTM` (crop label polygons in geographic and projected coordinate systems). Files with `wgs84` in their names use the

EPSG:4326 geographic coordinate system, while those with `utm` use the projected Universal Transverse Mercator system. Both systems are included for convenience: WGS84 for visualization tools like Google Earth, and UTM for precise analysis.

Shapefiles (`.shp`) contain vector data with coordinates; accompanying files `.shx`, `.dbf`, and `.prj` provide indexing, attribute data, and projection info. The `.dbf` files include fields such as `crop_name`, `apia_code`, `label`, `rgb_r`, `rgb_g`, `rgb_b`, and `hex_color` to describe crops and their colors.

Application scenarios and reference implementations using DACIA5 dataset

This section presents a set of use cases, proposed to illustrate the practical relevance, technical robustness, and methodological flexibility of the DACIA5 dataset. These representative scenarios demonstrate how the dataset can be effectively applied to real-world agricultural remote sensing tasks. By exploring these use cases, the objective is to showcase the versatility and value of the dataset, providing insights into how it can address real-world challenges.

Use case: Crop identification: past vs. present (Problem 1)

This task evaluates the ability to identify crops in the "present" (defined as early 2024) using only "past" data (2020–2023) for training. All 2024 data are reserved for testing, allowing us to assess how well past satellite observations generalize to unseen current conditions.

The DACIA5 dataset provides 32×32 patches extracted from Sentinel-1 SAR (2 channels) and Sentinel-2 MSI (12 channels) imagery. Patch extraction was conducted on parcels managed by the National Institute of Research and Development for Potato and Sugar Beet Braşov. The training dataset comprises 5431 Sentinel-2 patches and 5103 Sentinel-1 patches acquired between 2020 and 2023. The test set, based on 2024 acquisitions, includes 1023 Sentinel-2 and 892 Sentinel-1 patches. Temporal alignment between Sentinel-1 and Sentinel-2 was ensured by matching each multispectral acquisition with the closest SAR image.

Although the full dataset includes 17 crop types, only 12 are sufficiently represented to generate valid patches. Due to crop rotation and parcel size constraints, only 8 crop classes appear in the 2024 test set. Patch selection was performed to ensure dominant crop coverage and optimal phenological visibility.

Two classification experiments were designed: one using only Sentinel-2 data and another employing a fused format where each patch combines Sentinel-1 and Sentinel-2 channels (resulting in $32 \times 32 \times 14$ inputs), with SAR channels first, followed by multispectral bands.

Initial baseline experiments were performed using the ResNet18 architecture, chosen for its proven ability to handle complex visual patterns and its use of residual connections to alleviate vanishing gradient issues. This architecture is particularly suitable for classifi-

cation tasks involving small-sized inputs, such as the 32×32 patches in this dataset.

To adapt ResNet18 to the DACIA5 dataset, the initial convolutional layer was replaced with a randomly initialized layer to support the 12-channel input of Sentinel-2 data ($32 \times 32 \times 12$), instead of the standard 3-channel RGB input. In the fused experiment (Sentinel-1 AND Sentinel-2), the model accepted 14-channel inputs. The final fully connected layer was replaced with one comprising 12 output neurons, corresponding to the number of crop classes. Training was conducted using the AdamW optimizer with a learning rate of 10^{-3} and a weight decay of 10^{-3} for regularization. The loss function was cross-entropy. Batch size was set to 128. To address the dataset's inherent class imbalance, the *Imbalanced Dataset Sampler* from the *torch.utils.data* module was used. This sampling strategy ensured underrepresented crop classes were more frequently seen during training, reducing bias without requiring oversampling or augmentation.

Each experiment was repeated multiple times. For Sentinel-2-only data, test accuracy ranged between 60% and 65%, with an average of 62.67%. Training accuracy consistently exceeded 90%. In one representative run, the model achieved 91.72% accuracy on the training set and 62.13% on the test set.

For the fused Sentinel-1 and Sentinel-2 input ($32 \times 32 \times 14$), test accuracy was slightly lower, averaging around 55%, while training accuracy remained above 91%. One representative training instance reached 97.33% accuracy on the training set and 62.56% on the test set. The confusion matrices for these two cases are presented in Fig. 38 and Fig. 39.

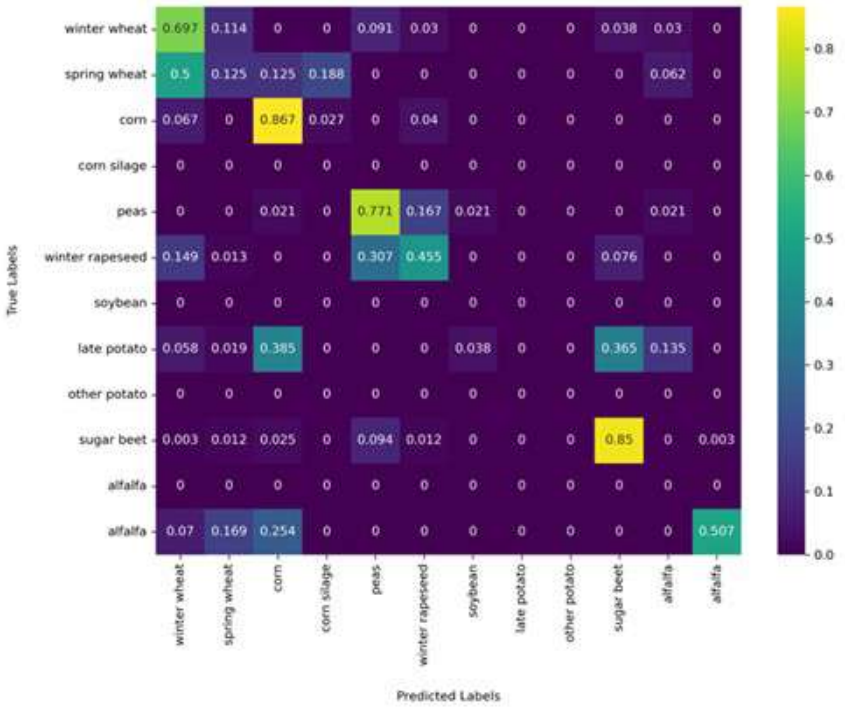

Figure 38: The confusion matrices obtained for two cases are detailed - for the Sentinel-2, the overall accuracy was 62.13%.

The classification results using Sentinel-2 alone and combined Sentinel-1 AND Sentinel-

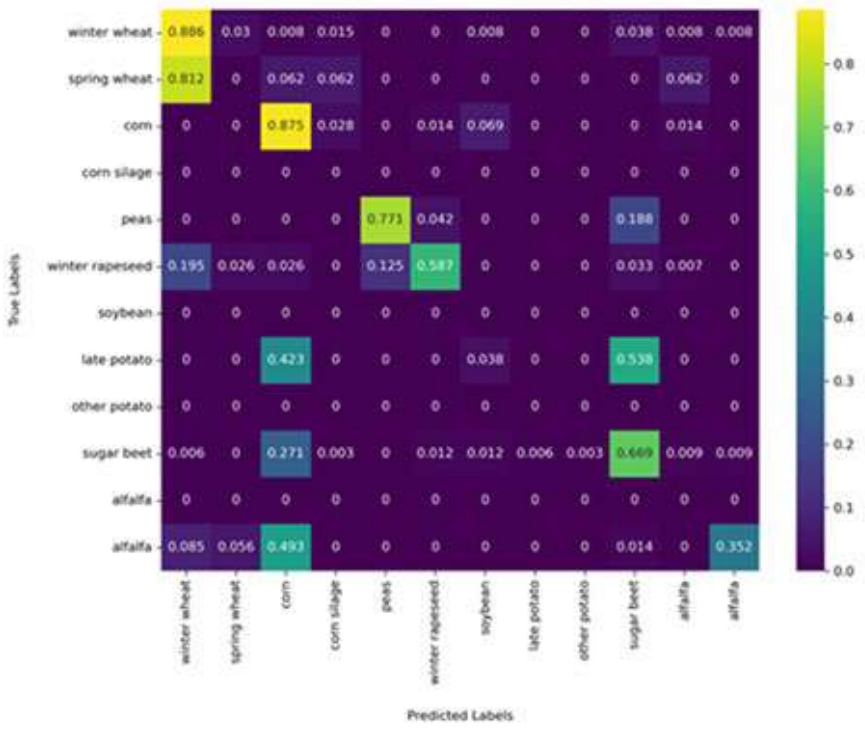

Figure 39: The confusion matrices obtained for two cases are detailed - for the fused Sentinel-1 and Sentinel-2 it was 62.56%.

2 data are broadly consistent, both in terms of accuracy and confusion matrix structure. This suggests that the model's overall performance is relatively stable regardless of the input source. However, only certain crops - such as corn and peas - are reliably identified (true positive rates ≥ 0.75), while others, like wheat and sugar beet, show moderate accuracy (true positive rates $\in [0.5, 0.7]$), and the remaining classes perform poorly.

Notably, the addition of Sentinel-1 data yields performance improvements for specific crops like wheat, potato, and rapeseed, due to radar sensitivity to crop structure and moisture. In contrast, for crops that are spectrally distinct but structurally similar, radar data may actually degrade classification performance.

As a complementary evaluation, an additional classification experiment was performed using a Random Forest (RF) classifier. In the first experiment, the model was trained using all Sentinel-2 patches (5,431 samples) of size $32 \times 32 \times 12$. In the second experiment, fused Sentinel-1 and Sentinel-2 patches (5,103 samples) of size $32 \times 32 \times 14$ were used. The test set included all corresponding patches from 2024. The model input consisted of multispectral pixels, each assigned the label of its originating patch. Comparing the results obtained by RF with those achieved by ResNet18 reveals certain similarities in classification, but also notable discrepancies. Some crop types are better identified by the deep learning model and less accurately by RF, and vice versa. Notably, crops such as winter wheat are consistently well separated by both approaches.

These two identification strategies serve as baseline references for subsequent benchmarking efforts. This use case further demonstrates the potential of the proposed dataset to support training of deep neural networks for agricultural crop classification using histor-

ical data - a task critical to precision agriculture, crop monitoring, and resource management. A detailed analysis and complete methodology are presented in [20].

Use case: Early crop identification (Problem 2)

The second experiment addresses early-season crop identification using carefully selected patches from DACIA5 dataset. The date of May 20th is considered a critical temporal threshold, aligning with the completion of sowing in Romania and the deadline for crop declarations set by APIA.

Using this temporal boundary, $32 \times 32 \times 12$ patches were extracted across 2020–2024, focusing on six major crops present throughout the growing season. This yielded:

- pre-May 20th: 1176 patches (young crops)
- post-May 20th: 1073 patches (mature crops)

Two distinct scenarios were investigated:

- Forward prediction: Training on early-season data (pre-May 20th) to identify mature crops (post-May 20th), simulating real-world crop monitoring needs
- Reverse learning: Training on mature crop data to identify young crops, exploring dataset robustness

Additionally, model performance was tested using merged Sentinel-1 and Sentinel-2 data (1162 pre-May patches, 1025 post-May patches) to evaluate whether multi-sensor fusion improves early identification accuracy. The same model architecture was maintained across all scenarios for consistent comparison.

Using ResNet18, the experiments using Sentinel-2 data showed that mature crop identification scenario achieved 41.01% test accuracy (91.79% training accuracy), while early crop identification reached 49.96% test accuracy (94.25% training accuracy). For the fused Sentinel-1 and Sentinel-2 data, mature crop identification yielded 33.01% test accuracy (90.31% training accuracy), while early crop identification achieved 41.16% test accuracy (95.17% training accuracy).

Through experiments and scenarios, several key patterns in model performance were observed. The results were notably influenced by class imbalance in the dataset, with wheat dominating the first scenario and alfalfa in the second. This imbalance is clearly reflected in the confusion matrices, where the model shows higher accuracy for majority classes.

Using Sentinel-2 data alone, the mature crop identification scenario demonstrated strong performance in detecting wheat, winter rapeseed, and peas. The early crop identification scenario showed different strengths, with superior accuracy in identifying alfalfa and peas. Notably, peas maintained consistent identification accuracy across both scenarios, suggesting robust spectral signatures throughout its growth cycle. The temporal

split point (May 20th) proved crucial for model performance, highlighting the importance of aligning dataset partitioning with crop phenological stages.

The integration of Sentinel-1 data with Sentinel-2 images did not yield substantial improvements in overall accuracy. The confusion matrices remained largely similar [20], with wheat and rapeseed maintaining high identification rates in the mature crop scenario, and alfalfa showing strong performance in early identification. However, a notable decrease in peas identification accuracy was observed with the fused dataset, suggesting that radar data might introduce noise rather than enhance discriminative features for this particular crop.

Through real-world use cases and experimental validation, the practical utility of the dataset in addressing critical agricultural challenges, particularly in early crop identification and monitoring, has been demonstrated. While acknowledging current limitations, the dataset's high-quality labels, combined with its potential for integration with other data sources, position it as a valuable resource for advancing precision agriculture research. As farming practices continue to evolve with technological advancement, datasets like DACIA5 play a crucial role in developing more accurate, efficient, and sustainable agricultural monitoring solutions.

The complete results presented in this section (II.2) have been published in *Big Earth Data* Journal [20], a Q1-ranked journal dedicated to Big Data applications in earth sciences and earth observation.

III. Insights and correlation studies in EO-Based agricultural monitoring

Understanding temporal dynamics and crop-specific relationships is essential for building adaptive and reliable AI models in precision agriculture. This chapter focuses on analytical studies that evaluate how machine learning pipelines perform under varying conditions across time and crop types.

The first study performs a multi-year analysis in the Brașov region, investigating interannual patterns and correlations between crop distributions and vegetation indices. The second contribution examines how data augmentation affects classification performance for different crop categories using Sentinel-2 data. Together, these studies provide empirical validation and strategic insights for optimizing model generalization in EO-based agricultural applications.

Highlighted studies and key findings:

- I.C. Plajer, **A. Băicoianu**, M. Debu, M. Ștefan, M. Ivanovici, C. Florea, A. Ghinea, L. Majercsik, 2025 - *Multi-year multi-crop correlation analysis in Brasov area* [134]

  Main findings:

  - Conducts a correlation analysis aimed at enhancing the interpretability of AI model performance in crop identification using NDVI time series.

  - The resulting correlation matrices revealed strong relationships not only for identical crops but also between different crop types, while in some cases, identical crops exhibited low correlation - suggesting potential misclassification risks.
  - Findings indicate that weather conditions significantly influence the correlation values, emphasizing the importance of environmental variability in temporal data analysis.
  - Experimental results using a FCNN aligned with the trends observed in the correlation analysis, highlighting both the strengths and limitations of using NDVI time series and ML models for crop classification tasks.

- L. Majercsik, **A. Băicoianu**, I.C. Plajer, 2025 - *Crop-Specific Effects of Data Augmentation on Classification Accuracy in Sentinel-2 Imagery* [106]

  Main findings:

  - The paper investigated the impact of data augmentation techniques on crop classification performance using Sentinel-2 multispectral imagery. The study reveals significant crop-specific variations in augmentation effectiveness, thus challenging common assumptions about universal data augmentation benefits in remote sensing applications.
  - The research used the DACIA5 dataset that includes five years (2020-2024) of Sentinel-2 imagery from the Brașov region of Romania, covering 12 different crop types with carefully validated in situ labels.
  - Peas exhibited consistently positive responses to most augmentations, with accuracy improvements of 10-25% across geometric transformations, suggesting robust distinctive spectral signatures that remain identifiable under substantial transformation. Winter rapeseed showed substantial benefits from rotation (up to +40%) and combined transformations (+10-20%), indicating that orientation variability improves classification accuracy for this crop. The paper highlights comparable findings for other crops as well.
  - Among individual techniques, random rectangle erasure emerged as the most balanced strategy, providing modest benefits for several crops with minimal negative impact on others. In contrast, random pixel dropout generally reduced accuracy for most crops, highlighting the importance of preserving pixel-level information for accurate classification.

III.1. Multi-crop correlation analysis in Braşov area

Remote sensing-based crop identification faces significant challenges in establishing reliable temporal patterns across multiple growing seasons. While NDVI time series from multispectral satellite data offer valuable insights into crop development, variations

in planting dates, weather conditions, and crop phenology complicate the creation of consistent classification models. Using the DACIA5 dataset, which comprises five years of Sentinel-2 multispectral data from 47 parcels with 17 crop types in Brașov North area, Romania, this section addresses the fundamental question of how temporal correlations in crop signatures affect machine learning-based classification performance. The dataset's unique advantage lies in its verified ground-truth labels provided by the National Institute of Research and Development for Potato and Sugar Beet from Braşov, enabling robust analysis of multi-year agricultural patterns and their impact on crop identification accuracy.

The study [134] presents temporal pattern analysis through spectral signatures derived from Sentinel-2 data across the growing season. Each crop type's development is represented through averaged spectral curves, capturing critical phases from planting to maturity. For example, Fig. 40 illustrates this methodology using late potatoes, where the reflectance curves highlight distinct growth stages. These temporal variations in spectral patterns serve as characteristic signatures for agricultural crop identification.

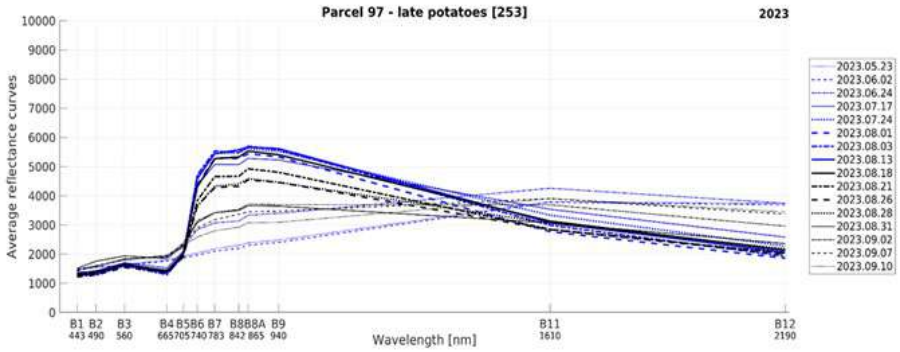

Figure 40: Temporal evolution of mean spectral reflectance - Late potato crop (parcel 97, 2023).

Potato crop development [113], [123] is commonly summarized into four main phenological stages: planting to sprouting, sprouting to bud flowering, bud flowering to flowering, and flowering to maturity. Sentinel-2 spectral data effectively captures these stages through specific band responses. During the initial post-planting phase, low reflectance is observed in vegetation-sensitive bands (B5-B7, B8, B8A, B9), indicating minimal chlorophyll activity. The growth development period shows gradual increase in reflectance values, particularly enhanced vegetation response in chlorophyll-sensitive Band 8. At maturity, peak reflectance indicates maximum chlorophyll production with dense vegetation signature across spectral bands. The senescence period is marked by declining reflectance patterns and distinctive changes in moisture-sensitive bands (B11, B12). This spectral evolution provides crucial indicators for crop monitoring, enabling precise tracking of growth stages and potential water stress conditions through moisture-sensitive bands.

Sentinel-2 multispectral images provide two essential spectral bands for NDVI calculation: B4 (664.6 nm, RED) and B8 (832.8 nm, NIR). Temporal NDVI analysis was performed across all parcels throughout the growing season, with Fig. 41 demonstrating the index evolution for sugar beet cultivation in parcel 39 during 2023. The analysis

encompasses thirty temporal observations at irregular intervals, determined by Sentinel-2 data availability and weather conditions. Sugar beet serves as a representative case study, given its regional significance and its importance in NIRDPSB research activities.

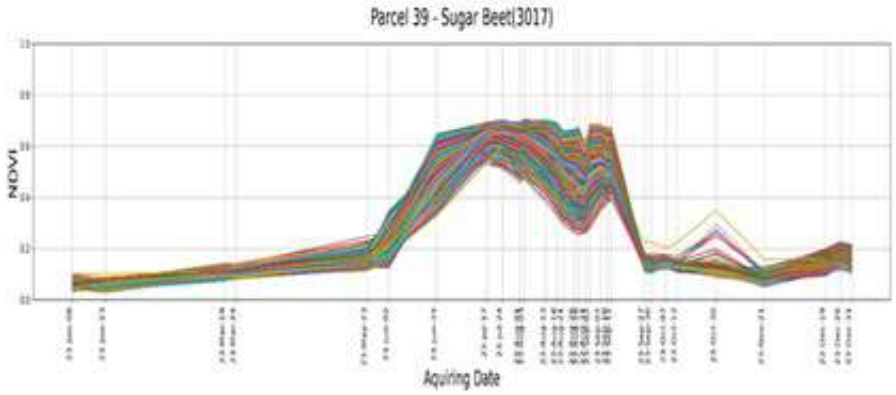

Figure 41: Temporal NDVI evolution - Sugar beet field analysis, parcel 39 (2023).

Representative NDVI curves were generated at two levels: parcel-specific and crop-specific. At the parcel level, NDVI values were averaged across all pixels within each individual parcel. Fig. 42 demonstrates this approach for parcel 39 in 2023, showing the averaged NDVI evolution (red curve). At the crop level, a comprehensive NDVI curve was calculated by averaging values across all parcels containing the same crop type. For sugar beet in 2023, this crop-specific curve (magenta) represents the average NDVI across three parcels, with seeding and harvesting events marked by blue arrows.

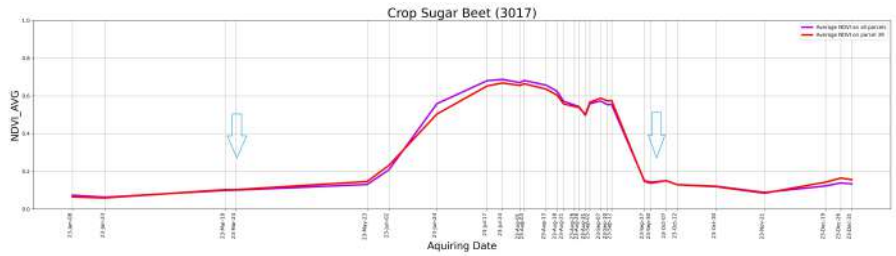

Figure 42: Comparative NDVI analysis. Individual parcel vs. crop-wide averages for sugar beet (2023).

Representative NDVI curves capture critical crop life cycle events, from growth through decay to harvest. In the case of sugar beet, peak NDVI values (≈ 0.7) marginally exceed the 0.66 threshold for healthy vegetation. This slight deviation from expected values may be attributed to multiple factors: satellite image processing artifacts, sensor aggregation effects, atmospheric conditions, and the characteristic sugar beet canopy structure which typically leaves soil partially exposed. NDVI time series analysis enables detection of growth anomalies within parcels by identifying deviations from representative patterns. Multi-year analysis of these curves, particularly when correlated with meteorological data, offers valuable insights into long-term crop health patterns.

The single-year correlation analysis employs Pearson correlation coefficients to evaluate relationships between representative NDVI time series across different crops in 2023.

Fig. 43 presents the correlation matrix of per-parcel average NDVI patterns, where labels indicate parcel indices and corresponding crop types. The correlation intensity is visualized through a yellow-to-green gradient, with darker green indicating stronger correlation between NDVI temporal signatures.

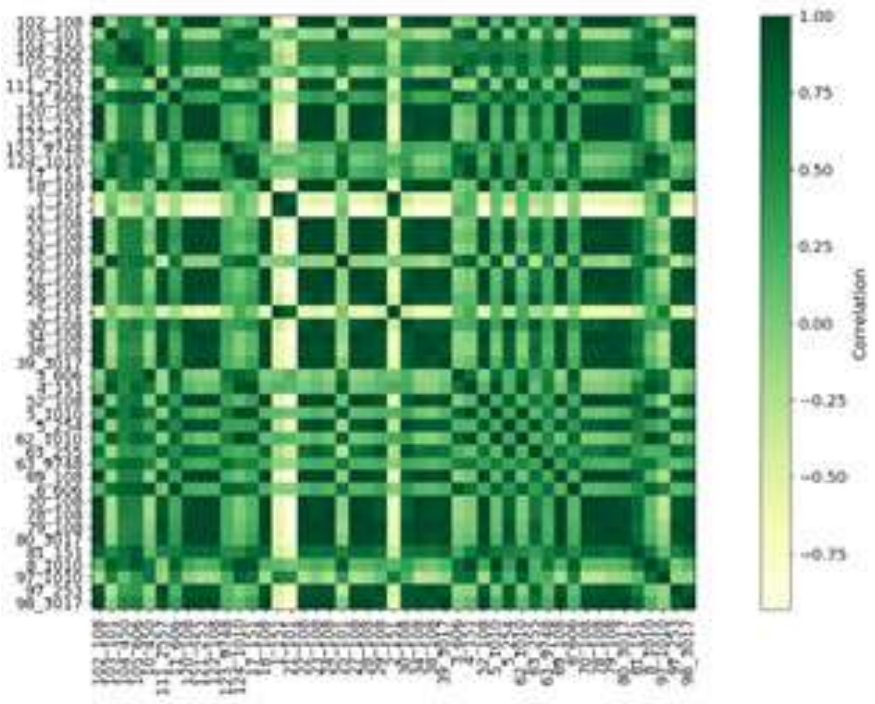

Figure 43: Correlation matrix of average NDVI time series per parcel.

Analysis of Fig. 43 reveals high correlation coefficients between parcels containing identical crops, as anticipated. Notable similarities appear between different cereal varieties (codes 101 and 108) and potato variants (codes 254 and 255), suggesting potential classification challenges due to their similar temporal NDVI patterns. Interestingly, certain distinct crops, such as sugar beet (parcel 39, code 3017) and corn (parcel 38, code 108), demonstrate unexpectedly high correlation despite their different biological characteristics.

The multi-year correlation analysis spans 2020-2024, with particular attention to weather variations, notably dry conditions in 2022 and 2024. The analysis addresses several methodological challenges stemming from data inconsistencies: variable crop presence across years due to rotation policies, different observation frequencies between years (30 dates in 2023 versus 40 in 2024), and irregular temporal sampling of Sentinel-2 data. To overcome these limitations, the methodology employed linear interpolation with 5-day fixed intervals, coupled with normalized cross-correlation to handle time series of different lengths. The correlation analysis utilized a shifting approach. These calculations were performed using the cross-correlation function provided by the SciPy library in Python. This function computes the correlation between two discrete sequences, x and y, representing the time series from two different years. Results visualization (Fig.44 and

Fig. 45) demonstrates the temporal alignment of NDVI patterns for key crops (corn-108, late potato-253, alfalfa-9748) between 2023 and adjacent years, accounting for seasonal shifts and environmental variations.

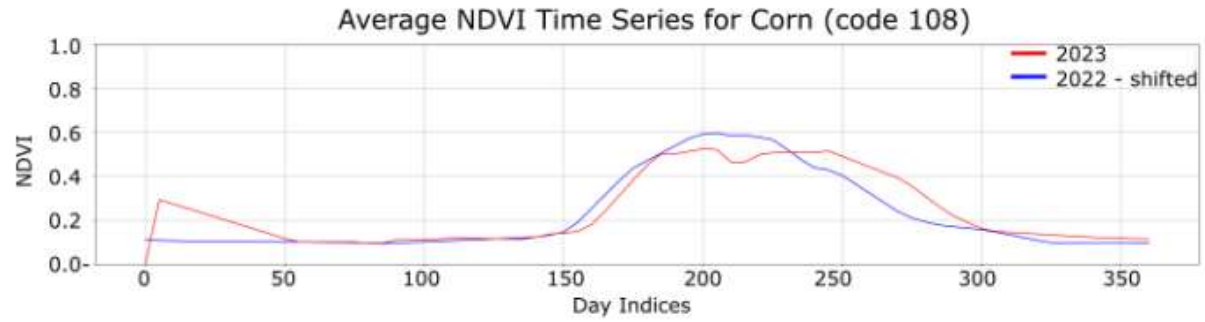

(a)

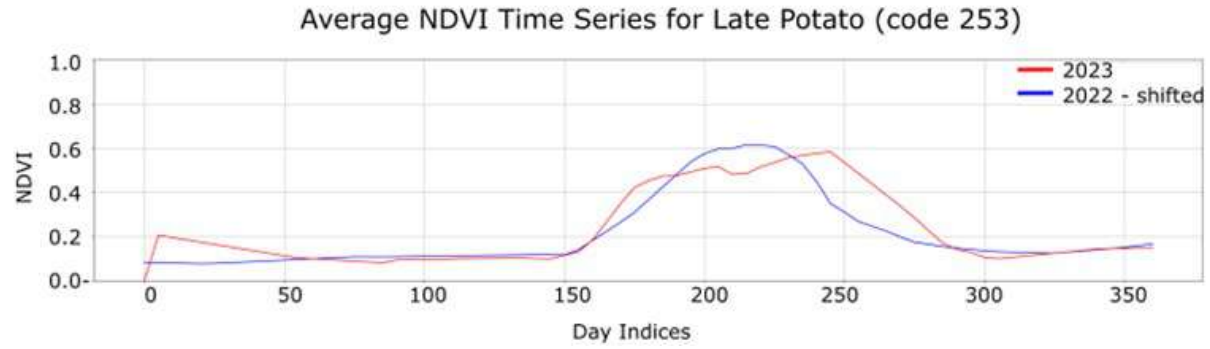

(b)

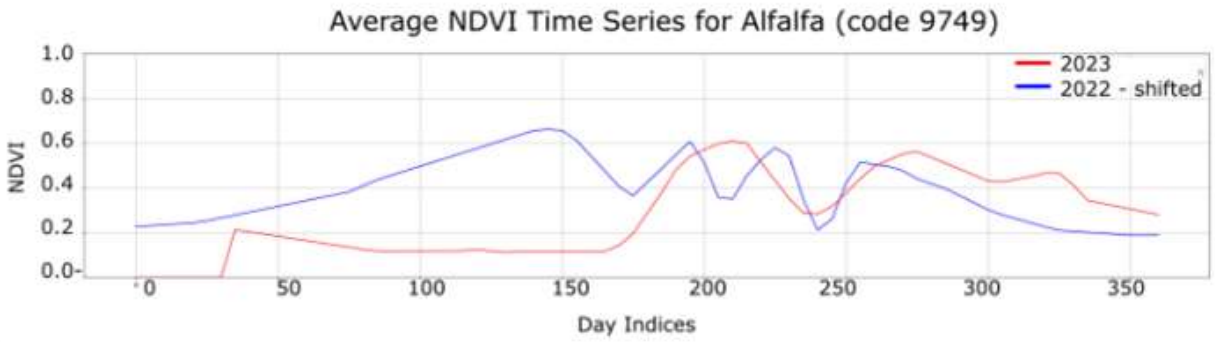

(c)

Figure 44: Pairs of interpolated average NDVI time curves for 2022 and 2023. The curves of 2022 are shifted, as to obtain the maximum correlation with the curve of 2023.

Cross-year correlation analysis was performed between consecutive pairs (2022-2023 and 2023-2024), with particular focus on weather pattern influences. The analysis is significant as 2022 and 2024 experienced similar dry conditions, while 2023 maintained normal to high precipitation levels. Using normalized cross-correlation of interpolated NDVI curves for common crops between year pairs, the analysis aimed to reveal whether weather similarities manifested in temporal crop signatures. Figure 46 presents these relationships through a heat map visualization of the cross-correlation matrix for 2022-2023, enabling direct comparison with single-year patterns observed in 2023.

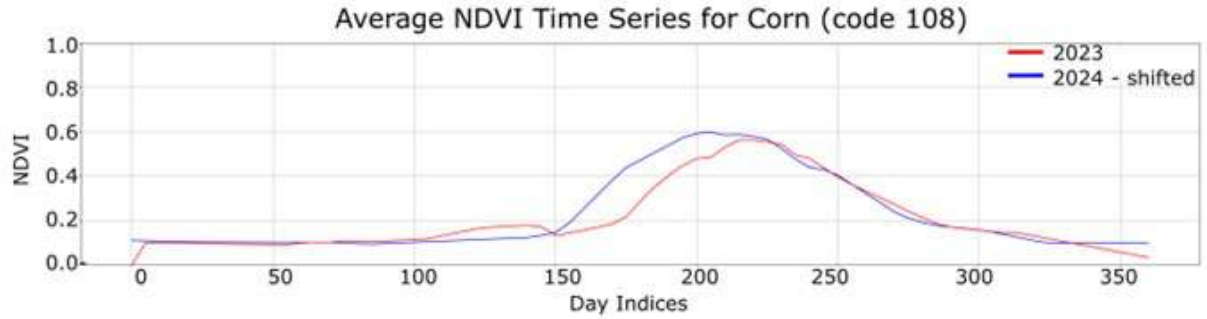

(a)

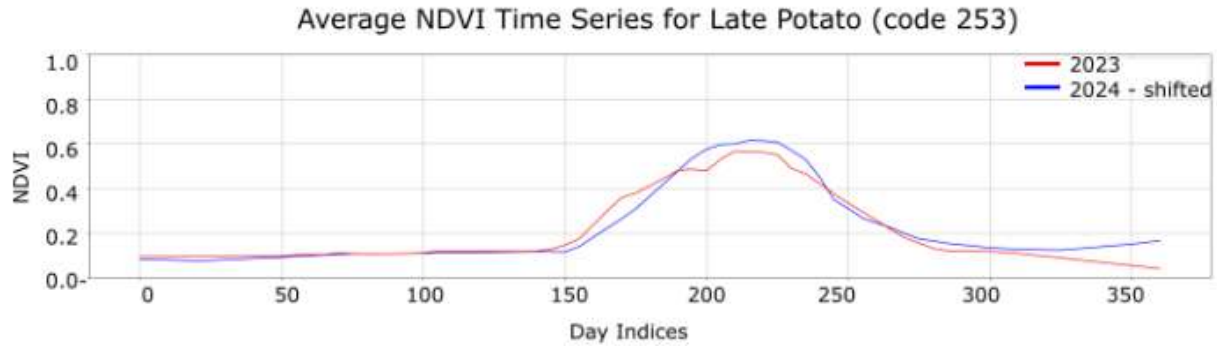

(b)

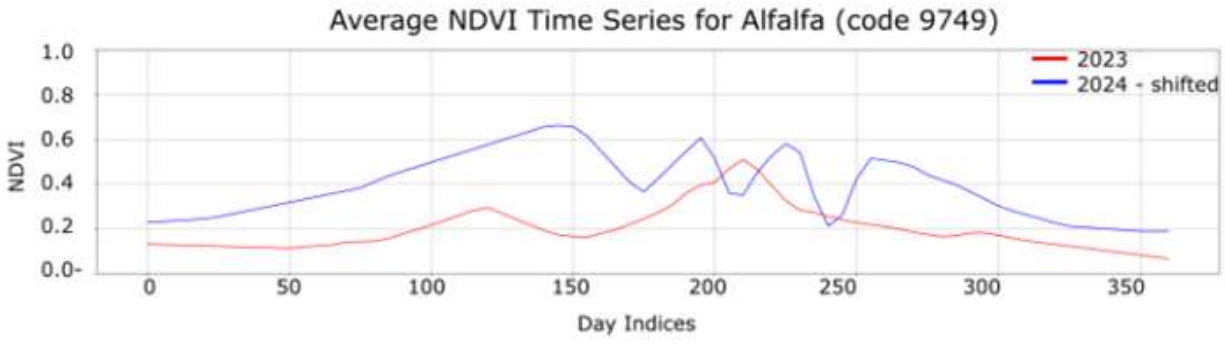

(c)

Figure 45: Pairs of interpolated average NDVI time curves for 2022 and 2023. The curves of 2022 are shifted, as to obtain the maximum correlation with the curve of 2023.

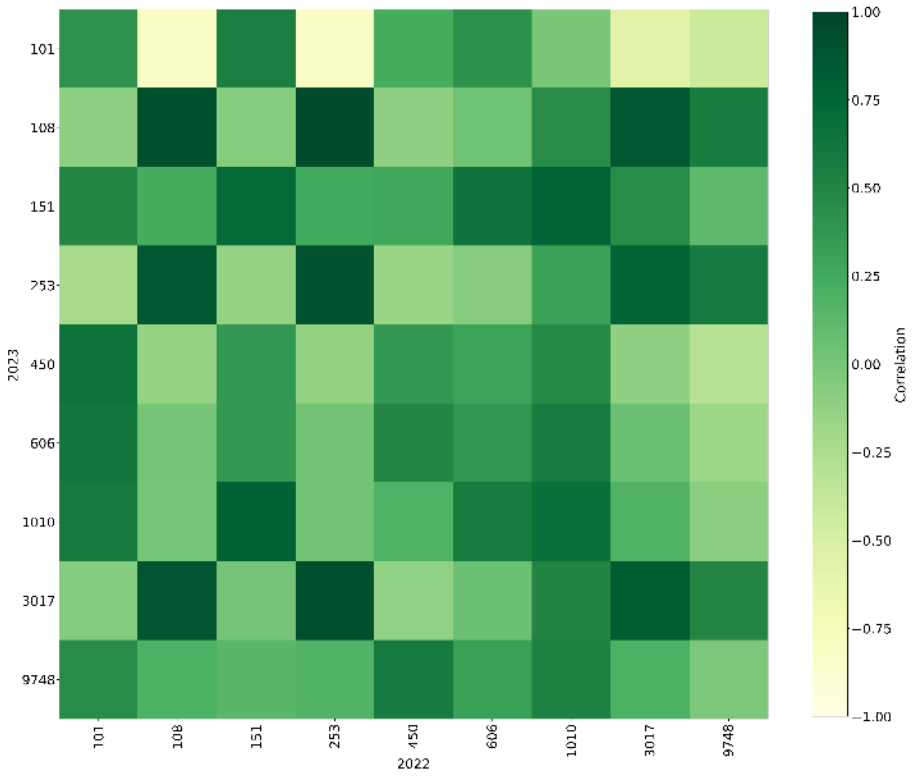

Figure 46: Cross-correlation matrix of average NDVI signature per crop. The pairs of interpolated signatures considered are from 2022 and 2023, where the 2022 signatures were shifted, as to obtain the best correlation.

To validate the correlation analysis findings, a machine learning classification approach was implemented using NDVI time series. The methodology employed a FCNN trained on interpolated data from 2022-2024. The training dataset comprised NDVI time series from 2023, with model validation performed on 2022 and 2024 data. All temporal sequences underwent interpolation over a complete annual cycle [0, 365], followed by downsampling to 5-day resolution, resulting in 73 values per sequence. Importantly, the 2022 and 2024 time series were temporally shifted to maximize correlation with their 2023 counterparts. While initial training showed promising results with 85% accuracy on the training set over 400 epochs, performance stabilization and potential degradation was observed after 300 epochs when testing on 2022 and 2023 datasets, indicating model overfitting. The complete experimental setup, classification results, and detailed performance analysis are thoroughly presented in [134], where the relationship between temporal correlations and crop identification accuracy is extensively examined.

III.2. Impact of data augmentation on crop-specific classification performance

This section examines the impact of data augmentation techniques on crop classification performance using Sentinel-2 multispectral imagery. Research presented in [106] revealed significant crop-specific variations in augmentation effectiveness, thus challenging common assumptions about universal data augmentation benefits in remote sensing applications.

The research used the DACIA5 dataset that included five years (2020-2024) of Sentinel-2 imagery from the Brașov region of Romania, covering 12 different crop types with carefully validated in situ labels. A temporal data split was implemented using 2020-2023 data for training (5,436 patches) and 2024 data for testing (1,017 patches), with only 8 of 12 crop classes appearing in the test set due to crop rotation practices. The study used a modified ResNet18 architecture adapted for 12-channel multispectral input, trained using an imbalanced dataset sampler to address significant class disparities in the data.

A total of seven augmentation strategies were subjected to systematic evaluation, namely random region swapping, random pixel dropout, random rectangle erasure, horizontal flip, vertical flip, rotation, and a combined approach combining all transformations. Each technique was designed to address specific aspects of satellite imagery variability, ranging from spectral variations due to environmental conditions to geometric transformations addressing positional biases.

These findings emphasise the need for customised augmentation approaches in crop classification systems. The study proposes crop-specific augmentation recommendations: geometric transformations for crops with symmetric structures, selective use of spatial obstruction/masking techniques for variable-response crops, rotation-based augmentation for crops with directional sensitivity, and minimal transformations for crops with strong orientation-specific features.

The analysis reveals that data augmentation effectiveness in crop classification exhibits strong crop-dependent patterns, challenging the conventional assumption of universal augmentation benefits. Based on experimental results illustrated across Fig. 47 and Fig. 48, specific augmentation strategies demonstrate varying impacts depending on crop characteristics. For instance, crops with symmetric structures respond positively to geometric transformations, while those with strong directional features benefit primarily from rotation-based augmentation. Notably, certain crops like sugar beet and alfalfa show performance degradation with standard augmentation techniques, suggesting the need for more nuanced approaches.

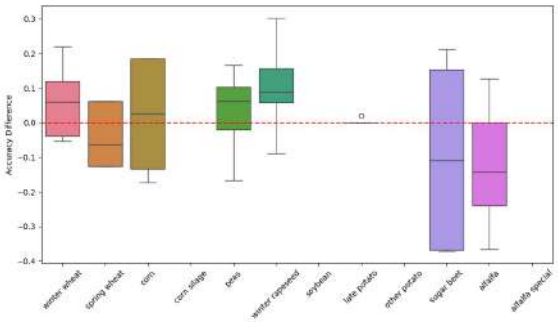

(a) Random Region Swap vs. the baseline

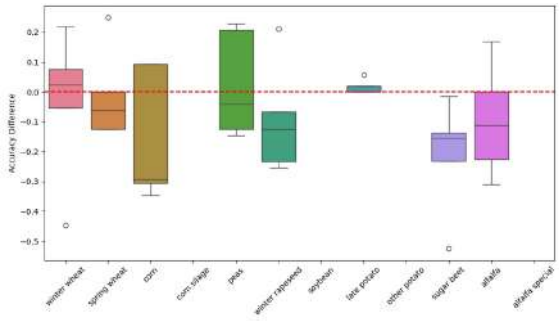

(b) Random Pixel Dropout vs. the baseline

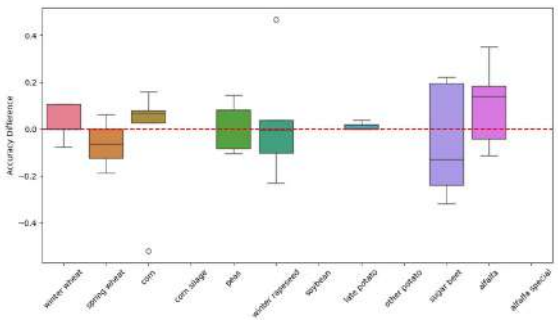

(c) Random Rectangle Erasure vs. the baseline

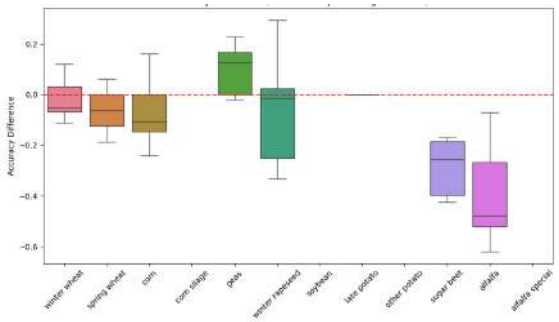

(d) Horizontal Flip vs. the baseline

Figure 47: Accuracy Differences (Part 1)

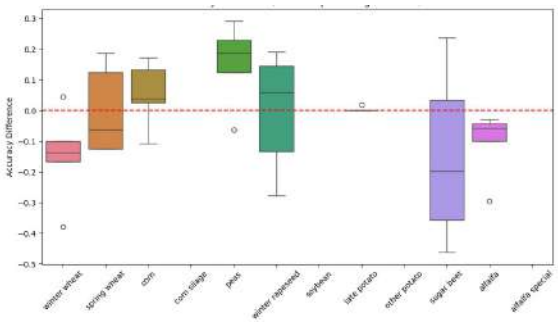

(a) Vertical Flip vs. the baseline

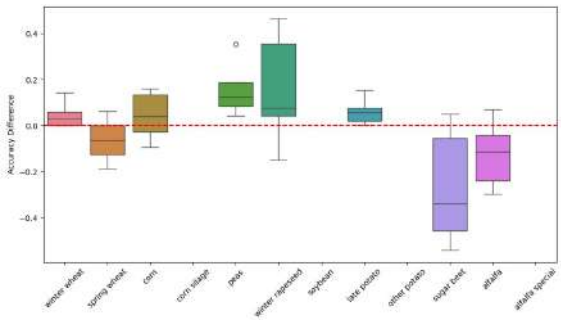

(b) Rotation vs. the baseline

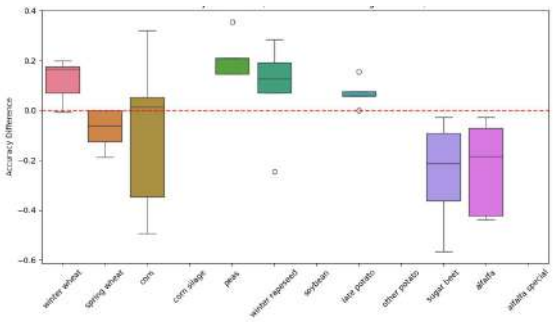

(c) All Transforms vs. the baseline

Figure 48: Accuracy Differences (Part 2)

The study, utilizing ResNet18 architecture on Sentinel-2 imagery, indicates that while some augmentation techniques enhance model generalization, persistent challenges remain in handling class imbalance and underrepresented crops. These findings emphasize

the importance of developing crop-specific augmentation strategies rather than applying universal transformations. Future research directions point toward exploring specialized architectures for multispectral data and developing advanced, domain-specific augmentation techniques to improve classification accuracy in precision agriculture applications.

The research presented in this section (III) is supported by two peer-reviewed publications in the field of Earth observation and remote sensing. The correlation analysis findings were presented at the *European Association of Remote Sensing Laboratories (EARSeL) Symposium* [134], a prestigious scientific network connecting European remote sensing institutes from both academic and commercial sectors. The data augmentation study was published in the proceedings of the *18th International Conference on Engineering of Modern Electric Systems 2025* [106]. These publications validate the methodological approach and contribute to the broader understanding of crop monitoring through Earth observation techniques.

Chapter 2. AI for monitoring, diagnosis, and optimization in complex systems

Problem statement and research context

This chapter explores the diverse applications of AI techniques for enhancing the performance, reliability, and adaptability of complex systems across multiple domains. From predictive maintenance in industrial environments to intelligent fault management in software systems, AI is transforming how we monitor and manage critical infrastructures.

The chapter explores five key areas where AI technologies have made significant contributions. The first part examines AI applications in condition monitoring and fault management, encompassing both industrial systems and software development. This subsection encompasses various case studies ranging from vehicle suspension and bearing failure detection to software bug management and big data processing solutions. The second subsection focuses on evolutionary algorithms, particularly genetic algorithms, and their application in structural system optimization, demonstrating their effectiveness in solving complex engineering problems. The third subsection explores AI's role in medical environments, specifically addressing emergency department triage optimization through machine learning approaches. The fourth subsection delves into enhancing time series prediction accuracy through novel data augmentation techniques. It details the application of fractal interpolation strategies in conjunction with LSTM models for generating high-quality synthetic data, offering valuable insights for applications in remote sensing and environmental monitoring. Finally, the fifth subsection investigates digital innovation in maritime applications, showcasing how AI-driven solutions are transforming marine energy management and facilitating technology transfer from automotive to maritime sectors.

The findings detailed in this chapter were developed within three major research projects. The *DITARTIS* project, through its research packages focusing on AI-based electrical machines and drives, alongside data-driven condition monitoring and predictive maintenance in EMS, provides the foundation for the industrial system optimization work. The *NEMOSHIP* project, which focuses on developing new modular electrical architectures and digital platforms for optimizing large battery systems on ships, supports the author's research in maritime applications. Both projects emphasize the critical role of AI in monitoring, diagnosis, and optimization of complex systems, directly supporting the research directions presented in this chapter. The *IMOCO4.E* project extends the author's research scope by developing distributed edge-to-cloud motion control intelligence for Human-in-the-Loop Cyber-Physical Systems, delivering AI and digital twin toolchains for resilient manufacturing applications, with particular focus on semiconductor, packaging, industrial robotics and healthcare domains. All three projects emphasize the critical role of AI in monitoring, diagnosis, and optimization of complex and interdisciplinary systems, directly supporting the research directions presented in this chapter.

I. AI for condition monitoring and fault management in industrial and software systems

Modern industrial and software systems demand increasingly high standards of safety, reliability, and efficiency throughout their operational lifecycle. Condition monitoring has emerged as a critical strategy in meeting these demands, offering multiple benefits: reduced maintenance costs, minimized unplanned downtime, sustained productivity, and enhanced operational safety. The integration of AI and ML technologies has revolutionized this field, enabling the processing and analysis of large-scale operational data for predictive fault detection.

Over the past decades, both academic research and industrial applications have demonstrated the significance of condition monitoring and fault diagnosis. The primary advantage of these systems extends beyond cost savings and efficiency improvements - they provide early detection of potential failures, thereby preventing catastrophic breakdowns and ensuring safe system operation [72]. By using AI-driven monitoring solutions, organizations can transition from reactive maintenance to predictive maintenance strategies, fundamentally transforming how complex systems are managed and maintained.

The following section presents a collection of research papers that address various aspects of condition monitoring and fault management across both industrial and software systems. These studies demonstrate the versatility and effectiveness of AI-based solutions in addressing complex monitoring challenges, from mechanical systems like vehicle suspensions and bearings to software development processes and big data management.

To support this research direction, the main findings from these studies, which implement and evaluate different AI-based monitoring and diagnostic approaches, are presented below. Each article contributes unique insights into the application of machine learning techniques for condition monitoring, fault detection, and system optimization. The research covers both traditional industrial applications (such as bearing failure detection and elevator monitoring) and modern software development challenges, such as bug triaging.

Highlighted studies and key findings:

- B. Hanganu, L.A. Radu, **A. Băicoianu**, 2020 - *Machine Learning for Condition Monitoring: Latest Trend and Review* [70]

  Main findings:

  - Successfully integrated multiple signal processing techniques (Fast Fourier Transform, Short Time Fourier Transform, and Mel-frequency Cepstral Coefficient (MFCC)) to create an enhanced feature extraction methodology from bearing vibration signals.
  - Applied Principal Component Analysis (PCA) for dimensionality reduction, enhancing the computational efficiency of the model training process.

  - Developed an innovative approach for transforming raw vibration signals into image representations suitable for deep learning analysis.
  - Designed and implemented a specialized CNN architecture for accurate bearing fault classification.
  - Demonstrated the model's capability to effectively distinguish between multiple fault states including inner race, outer race, and ball defects, as well as normal operation.
  - Demonstrated successful cross-dataset generalization by training on Case Western Reserve University data and validating on NASA Ames Prognostics Repository, achieving accurate fault detection despite different bearing specifications and operational conditions.
  - Proved MFCC's superior capability in extracting universal fault patterns, enabling accurate state classification across different datasets while effectively filtering out variable working conditions.

- **A. Băicoianu**, A. Mathe, 2021 - *Diagnose Bearing Failures With Machine Learning Models* [12]

  Main findings:

  - The experimental setup utilized the Spectra Quest Machinery Fault Simulator to collect comprehensive data across multiple sensors, including 3D accelerometers, acoustic emission sensors, and microphones, generating a substantial dataset of 262 million rows with 14 features for analyzing four distinct bearing conditions: healthy, inner race defect, outer race defect, and ball defect.
  - The Random Forest model demonstrated strong performance metrics, achieving 99.5% accuracy on training data and 82.4% on test data, with optimal hyperparameter configurations identified as `min_samples_split` at 5, `min_samples_leaf` at 100, and `n_estimators` at 100, though performance decreased when `n_estimators` exceeded 5000.
  - The data collection process encompassed a wide operational range, with speed measurements systematically varying from 300 to 2700 rpm in 120 rpm increments, while processing requirements varied significantly from under 2 minutes for random entry analysis to approximately 2 hours for full dataset processing.
  - The validation methodology incorporated both laboratory and real-world operational data, implementing cross-validation techniques and comparative analysis across multiple algorithms including k-neighbors, Random Forest, Multilayer Perceptron, AdaBoost, SVM, and Decision Tree to ensure robust model evaluation.

- A. Kerestely, **A. Băicoianu**, R. Bocu, 2021 - *A research study on running machine learning algorithms on Big Data with Spark* [86]

Main findings:

- An integrated data analytics system was developed, capable of fully processing large datasets generated during bearing manufacturing processes using a multinode cluster setup - a significant achievement compared to existing approaches that can handle only subsets of such data. Moreover, while the system performs equally well as other approaches on small datasets (up to 10 MB), it proves to be more efficient on larger datasets by employing a machine learning-based data processing core and a multinode cluster infrastructure.
- The comparative analysis revealed that while Scikit-learn performs better on smaller datasets, the Spark-based system showed superior scalability, achieving 5x performance improvement with proper multi-core configuration and demonstrating effective resource utilization across cluster nodes, particularly beneficial for large-scale industrial applications.
- The research produced a practical framework for implementing proactive fault diagnosis systems in manufacturing, with the system architecture designed to allow future optimizations and improvements, particularly in areas of accuracy enhancement and processing time optimization, making it valuable for real-world industrial applications.

- **A. Băicoianu**, P. Stănoiu, M. Velea, C. Husar, 2022 - *A Machine Learning Proposal for Condition Monitoring of Vehicle Suspension* [15]

  Main findings:

  - The study successfully developed and implemented a neural network approach for vehicle suspension monitoring, achieving prediction errors between 0-25% when comparing calculated versus predicted values, with particularly good performance in acceleration predictions showing nearly perfect overlap.
  - The research utilized Simcenter Amesim for comprehensive simulation, processing a substantial dataset of 1.05M records (reduced from 1.31GB raw data to 90MB processed data), with simulation parameters including varying body mass (375-450 kg), spring rates (23500-26500 N/m), and damper ratings (3150-6850 N/(m/s)).
  - The comparative analysis between classical machine learning approach and Simcenter Amesim's built-in neural network capabilities demonstrated that while Amesim offers automated and user-friendly features, the classical approach provided better flexibility and performance for complex suspension modeling.
  - The proposed methodology proved effective for predicting and preventing suspension-related defects, while also showing potential for broader applications including monitoring construction damage over time and improving simulation times for structural research.

- M.R. Raia, A. Ailincai, **A. Băicoianu**, C. Husar, C. Irimia, 2023 - *Condition Monitoring of Industrial Elevators Based on Machine Learning Models* [141]

  Main findings:

  - Successfully developed and implemented a Long Short-Term Memory (LSTM) model for real-time monitoring and fault identification in industrial elevator systems, providing an effective solution for condition state assessment.
  - Utilized Simcenter Amesim to develop comprehensive simulation scenarios analyzing rope sliding effects by varying friction coefficients (from 1.0, representing the ideal case, to 0.1) and cabin loads (from 0% to 100%). The simulations demonstrated how this phenomenon affects cabin trajectory control and overall system performance when the cable becomes partially or fully decoupled from the sheave.
  - The model achieved optimal performance through hyperparameter tuning, reaching 96.75% accuracy and 0.96 F1 score on the testing dataset, with Optuna framework proving essential in preventing overfitting by optimizing the LSTM layer parameters.

- R. Bocu, **A. Băicoianu**, A. Kerestely, 2023 - *An Extended Survey Concerning the Significance of Artificial Intelligence and Machine Learning Techniques for Bug Triage and Management* [27]

  Main findings:

  - The paper provides a comprehensive and constructive survey of existing AI/ML-based bug triage approaches, highlighting their advantages, limitations, and implementation contexts.
  - Key functionalities such as duplicate bug detection, bug prioritization, and developer assignment are shown to benefit significantly from AI-driven algorithmic models.
  - The paper formulates and answers three key research questions:

    RQ1: Is ai capable to effectively enhance and automate software bug triaging?

    The survey confirms that AI techniques, particularly ML models, offer accurate and computationally efficient solutions that improve the precision and speed of bug triaging.

    RQ2: identification of significant performance parameters and metrics.

    The study identifies and compares common evaluation criteria across the literature, including accuracy, precision, recall, F1-score, and computational overhead.

    RQ3: Open scientific problems and possible future developments to enhance software bug triaging processes.

Despite promising results, several open challenges persist — including model scalability, generalization to unseen projects, explainability, and integration into real-world CI/CD pipelines.

  - Each surveyed bug triaging approach is discussed in terms of its advantages and issues, both conceptually and in relation to real-world software development processes.

- D. Căliman, V. David, **A. Băicoianu**, 2023 - *An Analysis of Improving Bug Fixing in Software Development* [30]

  Main findings:

  - The study resulted in the development of a web-based tool, BugFix Suggester (BFS), designed to assist developers in the bug triaging process. BFS automatically suggests previously fixed bugs that are similar to newly reported issues. By extracting relevant keywords and applying text similarity techniques, the tool effectively retrieves historical bug reports that are most relevant to the current context. This approach enhances the quality of suggestions and significantly reduces the time required to identify and diagnose new bugs.
  - The solution demonstrated broad applicability across different industries through its standardized data format, seamless integration with existing bug tracking systems, and effective implementation of BERT algorithm for text analysis.
  - The modular system architecture, composed of a preprocessing core, a keyword extractor, a similarity engine, and a web interface, enables easy integration with external bug-tracking systems or Natural Language Processing (NLP) engines.
  - Experimental validation shows that the retrieved similar bugs are highly relevant, and the tool supports a more informed and efficient bug-fixing process in software development.

- **A. Băicoianu**, A. Kerestely, J.E. Ruiz Sarrio, J.A. Antonino-Daviu, 2025 - *A Practical Approach to Bearing Fault Detection using Pretrained ResNet18* [19]

  Main findings:

  - The pretrained ResNet-18 model achieved 95.96% validation accuracy in bearing fault classification through fine-tuning on time-frequency representations of vibration signals.
  - Demonstrated robust performance across both laboratory and real-world datasets, indicating strong generalization capabilities. Also, the approach proved particularly effective with limited labeled data while maintaining rapid training convergence.

- Successfully adapted the ResNet-18 architecture for industrial fault detection through effective processing of time-frequency representations of vibration signals, minimal computational requirements, and strong performance in variable operational conditions.
- Results validate the practical applicability of transfer learning for predictive maintenance in industrial environments, especially where labeled data is scarce.

All these research contributions have advanced the field of AI-based condition monitoring and fault management, with each study being disseminated through recognized academic channels, either at domain-specific conferences (ranked B or C) or in specialized journals (a Q2 category), as in the case of [27]. The diverse range of applications, from vehicle suspension systems to software bug management, demonstrates the versatility and effectiveness of AI-based approaches in both industrial and software domains.

In the following sections, in-depth analyses of two representative studies are presented, highlighting the main research trends in applying AI for monitoring, diagnosis, and optimization in complex and interdisciplinary systems. First, the development of machine learning models for industrial elevator condition monitoring is examined, demonstrating innovative approaches to real-time fault detection in mechanical systems. Second, the use of a pretrained ResNet-18 model for bearing fault detection is explored, showcasing the potential of transfer learning in industrial applications. These studies were selected as they represent significant advances in both the theoretical framework and practical implementation of AI-based condition monitoring systems.

I.1. Condition monitoring in industrial systems: The case of elevator diagnostics

Problem statement

Condition monitoring and fault diagnosis have become key areas of focus in both research and industrial applications. Their implementation plays an important role in enhancing system reliability by minimizing maintenance costs, avoiding unexpected downtimes, and sustaining overall productivity [72]. Beyond these operational and economic benefits, their greatest value lies in the early identification of emerging faults, which is essential for ensuring safe system operation and preventing potentially severe or irreversible damage.

Given that more than 12 million elevators operate daily worldwide, and their number and usage are rapidly increasing, maintaining them on a daily basis would require over half a million hours of manual labor, with a failure occurrence exceeding 164,000 times per day [109]. These faults result in increased operational expenses, and to ensure system safety and reliability, frequent inspections are necessary. In response, numerous studies have explored condition monitoring solutions. For instance, [126] proposes a remote condition monitoring approach using IoT devices to track elevator vibration and acoustic parameters, enabling remote data collection and fault detection. Similarly, [109]

presents the development of a cost-effective wireless sensor node for elevator monitoring, discussing the engineering challenges involved. In [56], a monitoring and warning system is introduced that uses cameras and smoke detectors to identify unsafe passenger behavior in elevators. Implementing such methods often involves the creation of a digital twin capable of tracking machine performance, predicting failures, and planning maintenance activities. Furthermore, the availability of large volumes of data generated by these systems makes the application of machine learning and artificial intelligence not only possible, but highly effective for extracting insights, predicting failures, and optimizing maintenance strategies [34].

This section presents a novel methodology for monitoring and identifying operational faults in industrial elevator systems. The primary objective is to develop a comprehensive approach for characterizing system operation modes and detecting abnormal operating conditions. The study specifically focuses on two critical fault scenarios: rope sliding on the sheave and sheave aging phenomena.

Individual machine learning models are developed for each fault type, which are subsequently integrated into a unified fault detection framework. To construct the fault identification models, this section describes the utilization of dynamic simulations performed on multiple elevator system configurations, where various operational parameters are systematically modified from a reference parametric model. The resulting simulation data is then processed through advanced machine learning algorithms, enabling the development of a robust fault detection system capable of identifying elevator operational anomalies. Additionally, the following sections present the validation process and performance metrics used to assess the system's fault detection capabilities.

System overview and fault analysis

An electro-mechanical elevator is powered by a motor that drives the sheave-rope system to move the passenger cabin. To reduce the motor torque required, the weight of the cabin and passengers is partially balanced by adding a counterweight. The electric motor used is a permanent magnet synchronous motor connected to an electronic unit, primarily represented by an inverter. The elevator also includes the cabin, which consists of a fixed mass representing the frame and a variable mass accounting for the different loads inside the cabin, as well as the cables and the control system that precisely manages the elevator's position. As the elevator moves, the cables change length, causing their stiffness, elasticity, and damping properties to vary depending on the cable length. Furthermore, the cable model considers a variable mass to account for viscous friction.

The graphical representation of the model developed in Simcenter Amesim can be identified in Fig. 49, with key system parameters detailed in Table 11.

When the elevator cabin is empty, the motor runs at a speed set by the controller, and the cabin moves mainly because the counterweight is going down. This causes a negative torque on the motor, and the cabin moves according to the set speed and conditions.

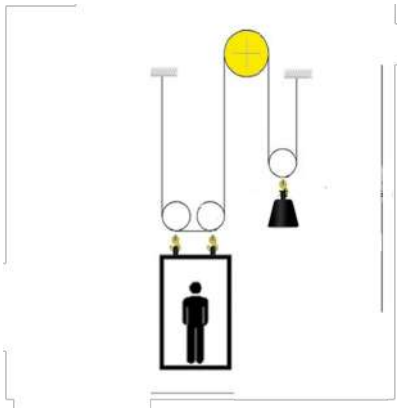

Figure 49: Elevator schematic representation.

Table 11: Technical specifications of the elevator

Parameter	Unit	Value
Cabin mass	kg	600
Counterweight mass	kg	1650
Sheave Radius	m	0.5
Gear ratio	-	1/10
Cabin viscous friction	N/m/s	1
Sheave viscous friction	Nm/rad/s	15
Rope Young's module	N/m^2	2.1×10^9
Maximum reachable height	m	75

The developed model simulates scenarios where rope sliding occurs on the sheave, characterized by changes in the friction coefficient between the cabin and rope. Fig. 50 illustrates the schematic representation of the rope sliding phenomena, where the rope (depicted in blue) interacts with the sheave (represented by the black circle).

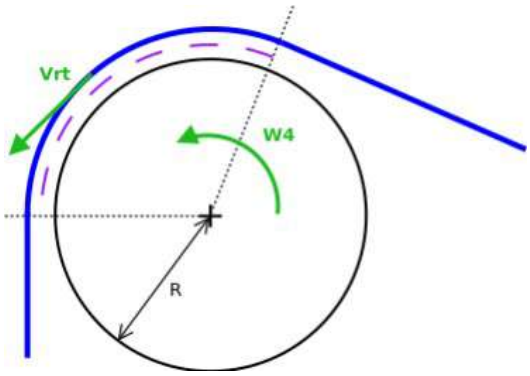

Figure 50: Sheave and rope.

During rope sliding events, the rope's movement becomes decoupled from the sheave's rotational movement due to changes in the dry friction force at their interface. This introduces a new variable, the tangential rope velocity (v_{rt}). This parameter is directly influenced by modifications in the dry friction coefficient and indicates whether the rope maintains ideal adherence to the sheave or experiences sliding conditions. The dry friction interaction between rope and sheave is represented by the purple dashed line in Fig. 50.

The capstan equation is employed to calculate the equivalent tangential force between the rope and sheave, expressed as in Equation (4).

$$F_{rope} = F_{resistance} e^{\mu\beta} \tag{4}$$

Equation (4) can be expressed as:

- μ represents the coefficient of friction
- β denotes the wrap angle
- $F_{resistance}$ is the smallest value of rope forces at both ends of the rope segment wrapped around the sheave

The resistance force is determined by the formula form Equation (5).

$$F_{resistance} = min(|F_{end1}|, |F_{end2}|) \tag{5}$$

The modeling of rope sliding effects in Simcenter Amesim involved varying the friction coefficient from 1.0 (representing ideal conditions) to 0.1. The study encompassed five distinct simulation scenarios, characterized by incremental changes in cabin nominal load from 0% to 100% in 25% steps:

- Scenario 1 (0% load): Empty elevator operation, considering only cabin mass and counterweight
- Scenario 5 (100% load): Maximum passenger capacity, incorporating total passenger mass, cabin frame mass, and counterweight

During rope sliding events, the partial or complete decoupling of the cable from the sheave significantly impacts system performance. This decoupling impairs the control unit's ability to:

1. maintain precise cabin positioning
2. ensure passenger comfort
3. meet performance targets

Under these conditions, the cabin's trajectory becomes primarily determined by the mass ratio between the counterweight and the total cabin load, rather than by the control system parameters.

The second fault analysis focuses on sheave aging, examining the degradation of both efficiency and diameter over time. Environmental factors such as rope friction, corrosive conditions, and humidity contribute to sheave deterioration, resulting in diameter reduction and altered torque characteristics.

The sheave's torque efficiency (e) varies based on the relationship between torque (T) and angular velocity (ω), following Equation (6).

$$T = \begin{cases} T_{sheave}e, \text{when } T\omega < 0 \text{ (braking)} \\ T_{sheave}/e, \text{when } T\omega > 0 \text{ (accelerating)} \end{cases} \quad (6)$$

A hyperbolic function manages the transition between these operational modes.

The Simcenter Amesim model incorporates two primary degradation factors:

- Sheave efficiency: varied from 1.0 (ideal) to 0.1, in 0.1 decrements
- Sheave diameter: reduced up to 14% from nominal value, in 1.25% increments

Similar to the rope sliding analysis, five load scenarios were simulated, varying the cabin's nominal load from 0% (empty cabin) to 100% (maximum passenger capacity) in 25% increments.

Methodology and experimental validation of intelligent elevator fault detection

The fault detection system employs a cascaded Long Short-Term Memory (LSTM) network architecture for processing multivariate time series data. As illustrated in Fig. 51, the model consists of two interconnected LSTM networks, where the second network utilizes the final hidden state of the first as its input. To prevent overfitting, dropout layers are implemented between LSTM layers. A linear layer followed by a SoftMax activation function completes the architecture, allowing it to classify inputs into three separate categories:

- Class 1: Non-faulty operation
- Class 2: Sheave-rope sliding condition
- Class 3: Sheave aging condition

The training stages can be synthesized by the following methodology:

1. Load the time series dataframes for each variable.
2. Compute 160 samples, where every sample is the i-th column from each dataframe.
3. Standardize samples using sklearn's StandardScaler.
4. Load samples into pytorch Dataset.
5. Random split the samples, 60% for training, 20% for validation and 20% for testing.
6. Train model using a batch-size equal to 96 for training and 32 for validating.
7. Evaluate classification metrics on the test dataset.

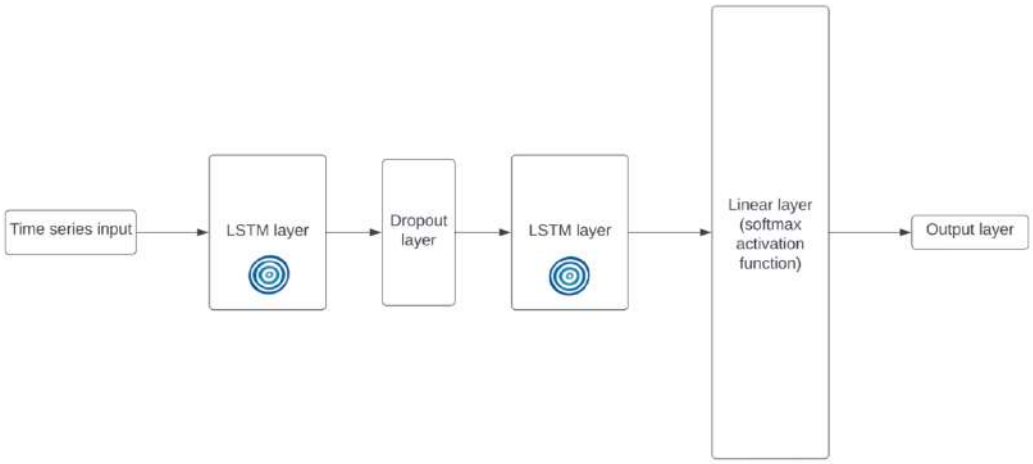

Figure 51: Neural network structure.

The dataset comprises 160 multivariate time series samples, each containing 6 variables over 70,000 time steps, namely break, cabin acceleration, cabin position, cabin velocity, motor speed, motor torque.

For every entry, the label is based on 2 constraints:

- For Class 1, the label is no defect
- For Class 2, the maximum value of the cabin position time series must be 4.01 meters.
- For Class 3, the condition for Class 1 must be satisfied and additionally the peak and mean value of the motor torque cannot be more than 20% higher/lower than the ideal case value.

After the labels for each entry are computed, the data can be standardized. Standardizing is a common scaling approach that converts the probability distribution for an input variable to a standard Gaussian by subtracting the mean from the values and dividing them by the standard deviation. Data standardization is performed using Sklearn's StandardScaler. The training implementation utilized an Adam optimizer and a Cross Entropy loss function, optimized for multi-class classification. Due to the limited size of the dataset (160 entries), batch sizes for the training and validation sets were specifically chosen to be equal to their respective sample counts, namely 96 and 32.

The model's performance over 200 training epochs is illustrated in Fig. 52, showing the evolution of accuracy, F1 Score, and validation/training loss.

The observed spikes in evaluation metrics can be attributed to the stochastic nature of the gradient descent algorithm and the dropout layer's neuron zeroing effect, a pattern also noted when implementing Gradient Clipping and L2 regularization techniques. The Tree-structured Parzen Estimator algorithm is employed for hyperparameter optimization, utilizing 500 trials to maximize both accuracy and F1 Score. The algorithm employs Gaussian Mixture Models to distinguish optimal parameter values through the ratio from

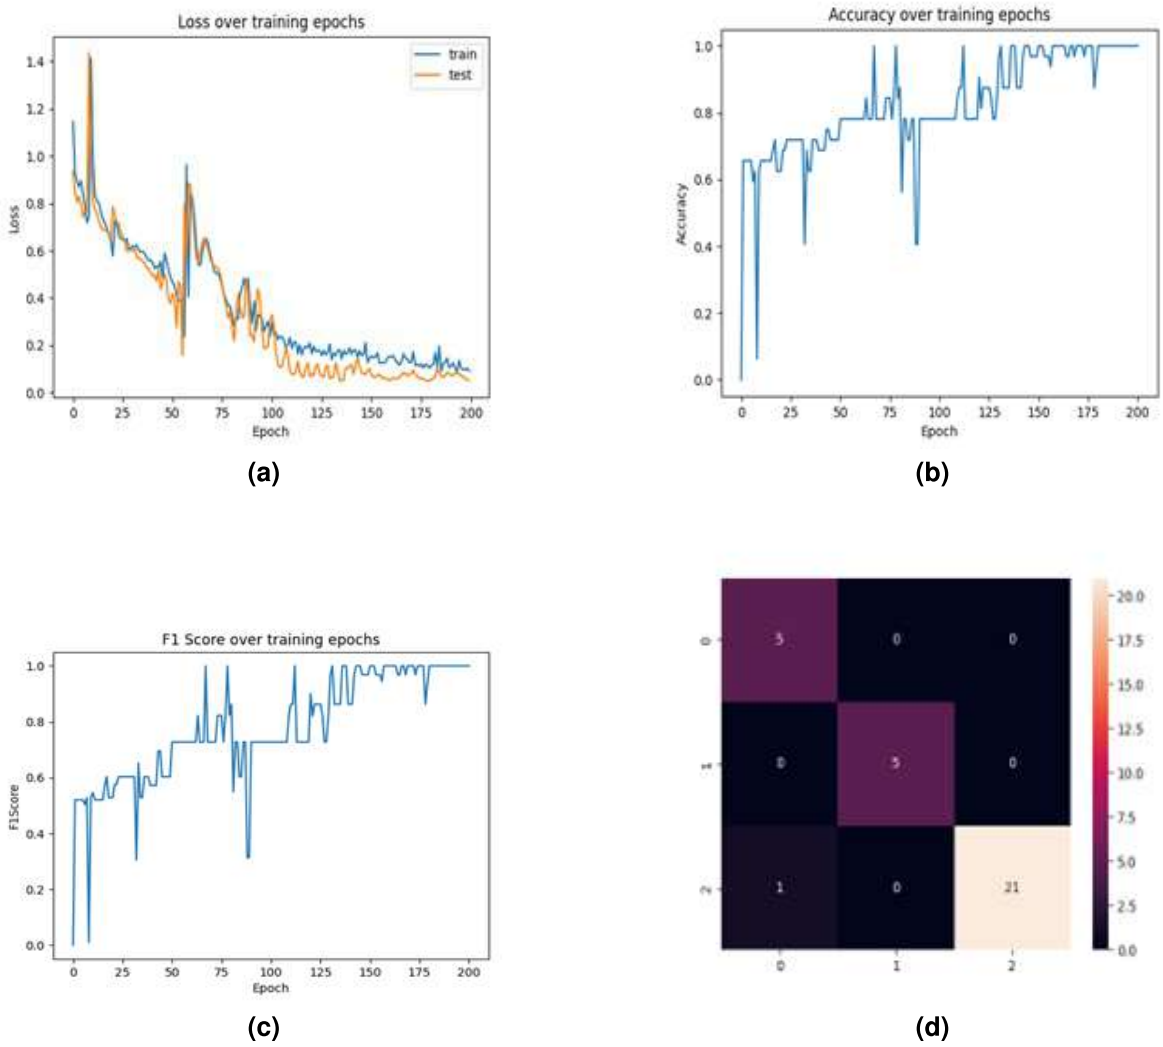

Figure 52: Loss, accuracy and F1 Score metrics and confusion matrix over the testing dataset.

Equation (7). In this $l(x)$ represents the density of optimal parameters (for the best hyperparameter values) and $g(x)$ the density of suboptimal ones (the other ones).

$$\text{Selection Criterion} = \frac{l(x)}{g(x)} \tag{7}$$

This implementation demonstrates effective fault detection capabilities while maintaining computational efficiency through careful architecture design and optimization strategies.

Several LSTM architectures were tested, ranging from 1 to 9 layers, learning rates spanning from 10^{-1} to 10^{-5}, dropout probabilities between 0.10 and 0.75, and hidden layer sizes varying from 32 to 512 units. Through systematic optimization, the model achieved optimal performance with a notably efficient configuration: a two-layer LSTM architecture, utilizing a 0.01 learning rate, 0.20 dropout probability, and 32 hidden units. This configuration demonstrated exceptional performance, achieving 96.75% accuracy and an F1 score of 0.96 on the testing dataset.

The effectiveness of `Optuna`'s optimization process is particularly noteworthy, as it prevented potential overfitting that could have resulted from excessive model complexity. Despite the wide parameter search space, particularly in terms of LSTM layers and hidden sizes, the optimization process identified that a more compact architecture could achieve superior performance while maintaining robust generalization capabilities.

Through systematic optimization and robust validation, this approach demonstrated exceptional performance in fault detection, achieving 96.75% accuracy and an F1 Score of 0.96 on the test dataset. These results strongly validate the effectiveness of integrating physical modeling with advanced machine learning techniques for elevator system monitoring. The developed framework provides a reliable foundation for predictive maintenance and safety enhancement in industrial elevator operations.

The complete methodology and results presented in this section (I.1) were initially developed as part of a research project and subsequently presented at the *IEEE International Conference on Emerging Technologies and Factory Automation (ETFA 2025)* [141]. This research was performed within the "Intelligent Motion Control under Industry4.E" (IMOCO4.E) project, which has received funding from the ECSEL Joint Undertaking under grant agreement No. 101007311.

I.2. Fault detection in industrial systems: The case of bearing diagnostics with deep learning

Problem statement

Bearing fault detection and classification represent major challenges in modern industrial maintenance, particularly in rotating machinery where bearings are consistently identified as the most failure-prone components [120]. The increasing complexity of industrial systems, together with the high costs associated with equipment downtime, has driven significant advancements in automated fault detection techniques [140], [84].

Recent developments in machine learning, particularly deep learning architectures, have transformed the approach to bearing fault diagnosis. While traditional methods relied heavily on manual feature engineering, modern deep learning techniques offer automated feature extraction capabilities, significantly improving the accuracy and reliability of fault detection systems [69]. Among these, CNNs have emerged as particularly effective tools, demonstrating exceptional performance in pattern recognition and adaptive feature learning [180], [153].

The availability of comprehensive datasets, such as those from Case Western Reserve University (CWRU) and Paderborn University (PU) [97], has enabled robust development and validation of these advanced diagnostic techniques. These resources have proven invaluable for benchmarking and improving fault detection algorithms, especially in scenarios involving variable operational conditions and limited labeled data.

This section expose an approach that explores bearing fault classification using a pretrained ResNet-18 model, fine-tuned on time-frequency representations of vibration

signals. By using transfer learning through a model pretrained on ImageNet, the common challenge of limited labeled data in industrial applications is addressed. The methodology demonstrates strong generalization capabilities, achieving a validation accuracy of 95.96% when evaluated on real-world vibration data [178]. This approach represents a practical solution for industrial fault diagnosis, combining the power of deep learning architectures with the constraints of real-world industrial applications.

Methodology and experimental validation of deep learning bearing diagnostics

The methodology is structured into three strategic stages, each addressing a critical component of the proposed bearing fault classification approach. The first stage focuses on transforming raw vibration signals into rich time-frequency representations, enabling the extraction of meaningful diagnostic patterns. The second stage introduces the adaptation of the ResNet-18 architecture, specifically engineered to handle the complexities of fault classification in rotating machinery. The final stage consolidates the approach through a robust transfer learning pipeline, where model training is fine-tuned to optimize performance on the target domain.

In industrial applications, rolling element bearings commonly develop faults on their critical contact surfaces, specifically the inner race, outer race, or the rolling elements themselves, which are typically spherical in radially loaded bearings. These faults typically manifest in two distinct patterns. The first pattern involves localized defects, such as spalls or pits, which create regular, periodic vibration patterns as rolling elements repeatedly impact the damaged areas. The second pattern consists of distributed defects that spread continuously along the bearing's circumferential surfaces, generating more complex and irregular vibration patterns. When bearings operate with these defects, the contact between rolling elements and damaged surfaces creates mechanical shocks that travel through the bearing structure as vibrations. These vibration signatures vary significantly based on several operational factors, including the bearing's physical dimensions, its operating speed, the location and type of fault, and how the damage progresses over time.

This approach enhances traditional envelope spectrum analysis by extending it into the time-frequency domain [146]. This improvement allows to track both amplitude changes and frequency patterns simultaneously, leading to more accurate fault detection. Time-domain vibration signals are converted into time-frequency representations to capture a more complete view of how bearing faults manifest. This enhanced analytical method proves particularly effective at distinguishing between different types of fault signatures: the regular impacts caused by localized defects, the irregular patterns from distributed damage, and combinations of multiple fault types. The specific frequencies associated with localized bearing faults, which are mathematically defined in [146], provide essential clues about both the type and severity of faults present in the system. This comprehensive approach to signal analysis significantly improves the ability to detect and classify

bearing faults compared to simpler, traditional methods.

Fig. 71 illustrate two distinct signals corresponding to inner race and outer race defects, respectively, each represented in terms of amplitude and frequency to highlight the characteristic differences in their spectral content. The time-frequency representation of such images indicates different signatures for different defects, which are utilized to fine-tune and validate the proposed network.

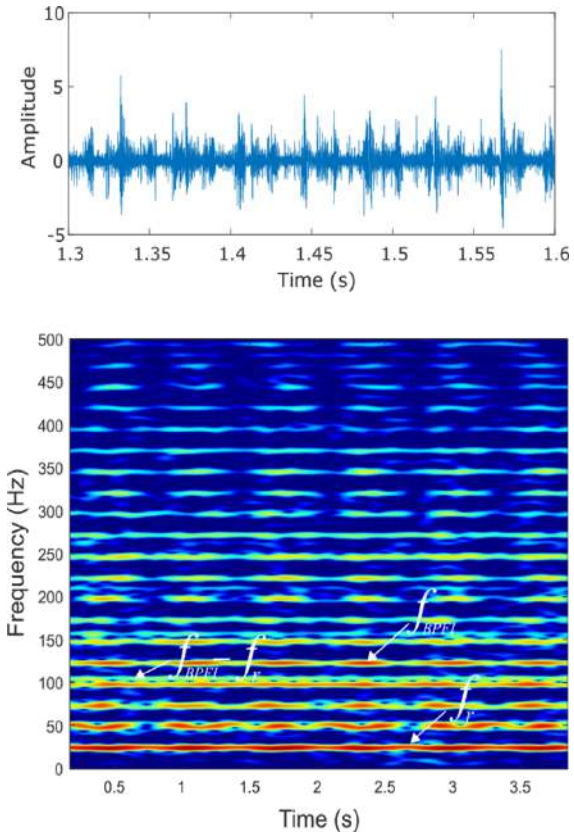

(a) Signals corresponding to inner race defects.

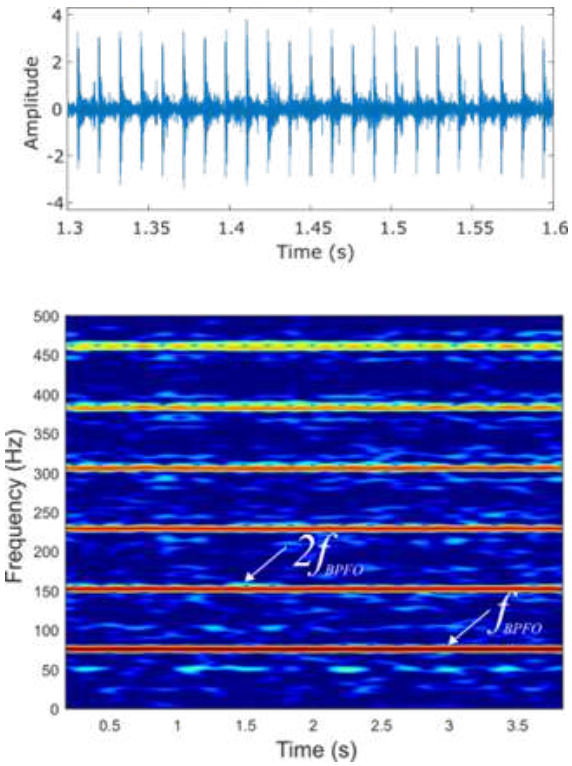

(b) Signals corresponding to outer race defects.

Figure 53: Illustrative examples of two distinct signals corresponding to inner race and outer race defects.

To address this problem, the ResNet-18 architecture was employed, a compact yet powerful convolutional neural network from the ResNet family, which revolutionized deep learning through its innovative residual learning approach. The architecture's distinctive feature lies in its skip/residual connections, which enable efficient gradient flow and mitigate the vanishing gradient problem common in deep networks. With approximately 11.7 million learnable parameters occupying only 43 MB of memory, ResNet-18 offers an optimal balance between computational efficiency and model capacity. Transfer learning is employed by using a ResNet-18 model pretrained on ImageNet, which is then fine-tuned for the specific bearing fault detection task. This approach is particularly advantageous in the industrial context, where labeled fault data may be limited. The pretrained weights provide a robust foundation of learned features that can be effectively adapted to recognize patterns in time-frequency representations of bearing vibration signals, significantly reducing training time while maintaining high classification accuracy.

The implementation of the proposed fault detection system uses a ResNet-18 architecture, with its classification layer adapted from the original 1,000-class output to a specialized three-class setup for identifying healthy bearings, inner race faults, and outer race faults. The transfer learning methodology follows a two-stage process designed to maximize the model's effectiveness in real-world applications, see Fig. 54.

Figure 54: Transfer learning workflow.

The first stage focuses on adapting the pretrained model using laboratory data collected under controlled conditions. During this phase, the network uses its pretrained ImageNet weights while learning to recognize specific patterns associated with bearing faults from artificially induced damage samples. This approach makes it possible to adapt feature extraction to the new task, while still keeping the strong visual recognition abilities learned from the original ImageNet training.

The second stage enhances the model's practical applicability through fine-tuning with real-world data obtained from life-accelerated tests. This crucial step addresses the inherent differences between laboratory conditions and actual industrial environments, capturing the variability and complexity present in real operational scenarios. By incorporating data from authentic bearing setups, the model develops improved generalization capabilities and enhanced reliability in practical applications.

The deployment strategy includes provisions for end-user customization, allowing further model refinement using site-specific data. This final adaptation step ensures optimal performance across diverse industrial settings, accounting for variations in operating conditions, equipment configurations, and sensor placements. This comprehensive approach creates a robust and adaptable fault detection system capable of maintaining high accuracy across different industrial environments while minimizing the need for extensive retraining.

This transfer learning approach demonstrated robust performance in real-world fault detection scenarios, achieving a validation accuracy of 95.96% and an F1 score of 0.9553. The confusion matrix analysis, see Fig. 55 reveals particularly strong performance in identifying inner race defects, with perfect classification accuracy for this fault category. However, the analysis also identified specific classification patterns where some outer race faults were occasionally misidentified as either inner race defects or healthy conditions, and certain healthy samples were incorrectly classified as having inner race defects.

The experimental configuration, including data distribution across training stages and optimization parameters, is detailed in Table 12.

The implementation of the pretrained ResNet-18 architecture proved particularly advantageous in the industrial context. The model's extensive pretraining on ImageNet's

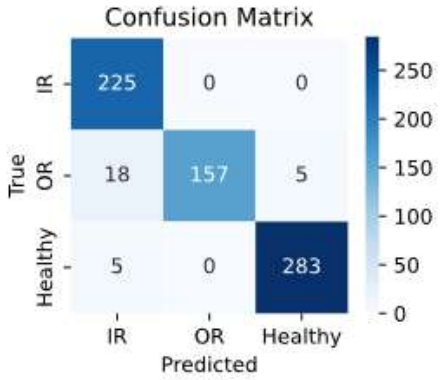

Figure 55: The confusion matrix, where IR is inner race and OR is outer race.

Table 12: Experimental setup summary

Parameter	Value
Real-world data (Fine-tuning)	25%
Real-world data (Validation)	75%
Healthy data (Stage 1)	10%
Healthy data (Stage 2)	10%
Healthy data (Validation)	80%
Learning rate	0.001
Momentum	0.9
Optimizer	SGD

diverse dataset of 1,000 classes provided robust feature extraction capabilities that transferred effectively to bearing fault detection. The residual architecture's inherent ability to learn deep hierarchical features contributed significantly to the model's strong generalization performance, even with limited training data. These characteristics, combined with the model's efficient fine-tuning capabilities, establish ResNet-18 as a practical and reliable solution for industrial bearing fault classification applications.

The complete methodology and results presented in this section (I.2) were initially developed as part of a research project and subsequently presented at the *International Conference on Future Energy Systems (FES 2025)* [19] in the track of *AI and Advanced Analytics in Energy Systems*. This work was supported by the European Union's Horizon Europe research and innovation programme, HORIZON-WIDERA-2021-ACCESS-03-01, under grant agreement No. 101079242 (DITARTIS project).

II. Evolutionary algorithms for structural system optimization

In any engineering or manufacturing context, a core objective is consistently to minimize resource consumption while ensuring optimal performance. Whether we're talking about packaging a product or constructing a complex building, the goal is often the same: use the least amount of material necessary to achieve the desired functionality and structural integrity. This fundamental optimization challenge, directly impacting cost, sustainability, and overall project efficiency, is precisely what is addressed here. Our focus is on using genetic algorithms, a powerful class of evolutionary computation methods, to significantly optimize the material quantity used in structural designs.

Currently, a substantial portion of structural design and optimization still relies on

highly iterative, manual processes. Engineers and architects often find themselves in repetitive cycles of design, calculation, and refinement, exploring a limited subset of potential solutions to arrive at something "good enough". This hands-on approach is time-consuming, prone to human error, and frequently results in designs that are compromises rather than true optima. The inherent difficulty of simultaneously achieving robust structural performance and cost-effectiveness through purely manual methods creates significant bottlenecks, extending project timelines and increasing the risk of design flaws.

Over time, various applications have been developed to streamline the work involved in designing and optimizing structures. However, most of these applications are primarily rooted in architectural design, aiming to create visually unprecedented structures. While they offer significant benefits in terms of 3D visualization of the final result, it prompts the question of whether aesthetic appeal, achieved at the cost of massive material waste, is truly sufficient. Furthermore, despite their advantages, the high costs associated with many of these applications can pose a significant impediment for smaller companies, often leading them to revert to older, manual working methods that involve dozens or even hundreds of lines of mathematical calculations, exponentially increasing the potential for errors and project delays.

Several notable software solutions exist in the field:

- Allplan is recognized as one of the most powerful Building Information Modeling (BIM) applications, providing capabilities for high-quality sketch creation. A major drawback, however, is the absence of a mechanism for monitoring, evaluating, and interpreting the designed constructions. This makes it challenging for engineers to identify what is lacking to achieve the initially proposed objectives.

- ACE OCP (American Computers and Engineers - Optimization Computing Platform) is positioned as an early commercial structural optimization tool, ACE OCP's primary limitation is its "surface-level" optimization approach. This implies it might not thoroughly explore the design space, potentially settling for local optima rather than achieving the maximum possible material efficiency.

- OptiStructure offers structural optimization services utilizing AI techniques. Similar to most existing frameworks, its optimization is performed under aesthetic constraints, ultimately yielding an imposing structure that still has room for further optimization.

- Altair employs advanced Finite Element Methods to analyze structural behavior under various loads. While excellent at validating integrity and predicting performance, it functions primarily as an analysis tool rather than an autonomous generative optimization engine. Notably, Siemens acquired Altair Engineering Inc. in March 2025, strengthening its simulation software portfolio.

This is precisely where modern advancements in AI, particularly in the domain of evolutionary algorithms like genetic algorithms, offer a transformative solution. These

techniques have matured significantly and are now powerful enough to tackle highly complex, multi-objective optimization problems. Unlike traditional methods that analyze a single design or iterate through limited permutations, genetic algorithms can actively explore vast, non-linear design spaces. By simulating natural selection processes, they can "evolve" designs that effectively balance multiple performance criteria: such as structural integrity, cost, and crucially, material efficiency, in ways that are simply unattainable through manual calculations or conventional software approaches. It's about empowering the computational system to discover the most efficient path through an almost infinite number of design possibilities.

To concretely illustrate the power and applicability of evolutionary algorithms in tackling this complex engineering problem, this chapter will explore two distinct case studies developed through collaborative research efforts. These studies directly address the aforementioned challenge of structural material optimization, demonstrating practical implementations of genetic algorithms.

The first study [16] offers a broader application of these principles to achieve enhanced material efficiency and performance in various structural contexts. The second [38] focuses on exploring the design space through parametric variations, allowing for the discovery of highly efficient structural configurations. Both works showcase how genetic algorithms can be effectively applied to derive innovative and resource-efficient structural solutions, moving beyond the limitations of traditional design methodologies.

These studies bring forward the following key elements relevant to the topic of this section:

- **A. Băicoianu**, A. Garofide, R.I. Luca, M. Vlădărean, 2022 - *Structural Optimization Using Genetic Algorithms* [16]

  Main findings:

  - Validated the efficacy of GAs in addressing complex, real-world structural optimization problems, demonstrating their practical applicability beyond theoretical exercises.
  - The study proposed a fully automated solution for the structural optimization process.
  - The open-source Project Chrono C++ library was used for performing robust Finite Element Analysis (FEA), which was instrumental in accurately simulating structural resistance and behavior under diverse force conditions.
  - Identified the formulation of a robust fitness function as a primary methodological challenge, requiring extensive iterative testing. This process was critical for effectively balancing competing objectives, such as material reduction (quantified by removed finite elements) and structural integrity (maximum stress), by assigning appropriate weighting to each variable.

  - Employed multiple result analysis techniques, including plotting the genetic algorithm's population evolution and visualizing the optimized structure using the Irrlicht library integrated with the Project Chrono framework.

- A. Cherdivar, **A. Băicoianu**, I. Popa, 2024 - *Parametric structural optimization using genetic algorithms* [37]

  Main findings:

  - Demonstrated the practical effectiveness of GAs in solving real-world structural optimization challenges.
  - Proposed a fully automated solution for structural optimization, offering a robust alternative to traditional engineering methods like MATLAB's Optimization Toolbox `fmincon` function.
  - Addressed two distinct practical engineering problems:
    * Determining the force on a beam to assess the need for adding a truss.
    * Calculating the minimum material required for a trussed beam to withstand a specific force.
  - Conducted an in-depth comparative analysis of various genetic algorithm configurations, evaluating performance across different selection methods (Roulette, Truncation, Tournament, Boltzmann, Stochastic Universal Sampling), crossover techniques (single-point, two-point, uniform), and mutation strategies. This analysis also included testing two distinct algorithm execution configurations under various iteration schemes.

II.1. Structural optimization using genetic algorithms

Problem statement

In structural engineering and design, a persistent and critical challenge is the optimization of material usage while simultaneously ensuring structural integrity, performance, and cost-effectiveness. Traditional manual calculation methods, though widely employed, are inherently iterative, time-consuming, and highly susceptible to human error. Recent advancements in AI, particularly in the domain of evolutionary algorithms like Genetic Algorithms (GAs) [93], offer a promising paradigm to overcome these limitations. Building upon this, the core problem addressed by this section is the development and implementation of a GA capable of optimizing a given initial structure under various engineering constraints, including material properties, applied forces, and stress limits.

The application of GAs in structural optimization has been extensively explored across multiple research efforts [165], [1], [90]. These studies often differentiate themselves through variations in gene encoding, selection and crossover methods, and the formulation of fitness functions, each tailored to achieve optimal solutions for specific structural

components and constraints. While GAs inherently provide heuristic, near-optimal solutions rather than guaranteed global optima, the poposed approach emphasizes rapid convergence and practical applicability. Furthermore, the integration of graphical visualization for intermediate and final designs will offer intuitive feedback, allowing users to assess results and refine parameters effectively. Ultimately, this section's output seeks to demonstrate the practical utility of GA in real-world structural optimization problems [89], offering a more efficient and automated solution that allows human experts to focus on tasks requiring unique human judgment and creativity.

Evolutionary algorithm architecture

The proposed GA is designed to simulate the principles of natural selection to identify optimal structural solutions. Inspired by Darwinian evolution, the algorithm iteratively refines a population of candidate designs, allowing only the "fittest" individuals to contribute to subsequent generations. This process involves two primary mechanisms: the selection of superior individuals to propagate, and the application of genetic operators (crossover and mutation) to generate diverse offspring. This GA is specifically adapted for structural optimization, aiming to produce designs that respond efficiently to constraints and applied forces.

A critical component of this evolutionary process is the *fitness function*, which quantifies the "suitability" of each candidate solution. While deriving an optimal fitness function often involves empirical "trial and error" extensive testing led us to the formula from Equation (8).

$$F = R^2 \cdot (max_S - max_{SS}) \tag{8}$$

Here, F represents the fitness value, R denotes the number of removed elements from the initial structure (indicating material reduction), max_S is the maximum allowable stress within the structure, and max_{SS} is the actual maximum stress experienced by the structure. This function rewards designs that achieve significant material reduction while maintaining stress levels below the defined limits.

The evolutionary cycle begins with the selection step, where individuals from the current generation are chosen to become parents for the next. The roulette wheel selection method [187] is employed, a stochastic approach where an individual's probability of selection is directly proportional to its fitness value. This ensures that chromosomes with higher fitness are more likely to be chosen, though even less fit individuals retain a small chance of being selected, contributing to population diversity. For each individual, a cumulative probability of selection is computed, guiding the selection process for the subsequent crossover step.

Following selection, the crossover step combines genetic material from the selected parents to produce offspring. Our implementation utilizes *single-point crossover*, where parts of the parents' genes are interchanged. In the context of the structural optimization problem, this operation directly translates to adding or removing material elements from

the structures represented by the individuals, effectively creating new design permutations.

Finally, the mutation step introduces random alterations to a limited number of genes within the offspring population. This operator is crucial for diversifying the new generations and mitigating the risk of premature convergence to local optima. For this problem, mutation involves either adding or removing elements from the structural mesh. Given that the mesh elements are represented by binary values (0 indicating the absence of a structural element, 1 indicating its presence), we use bit-flip mutation, which simply inverts the value of a selected gene. This mechanism ensures continuous exploration of the design space, preventing the algorithm from getting stuck in suboptimal regions.

The overall flow of the proposed GA is in Fig. 56. The process initiates with population initialization, where, instead of random generation, pre-built structures are employed to guide the algorithm toward maintaining desired empty spaces. Each individual's fitness is then computed, followed by the iterative application of selection, crossover, and mutation. A consistent binary encoding is utilized, as illustrated in Fig. 57. This cycle repeats until a termination condition is satisfied (e.g., a set number of generations or a sufficiently optimized solution), resulting in a final optimized structure.

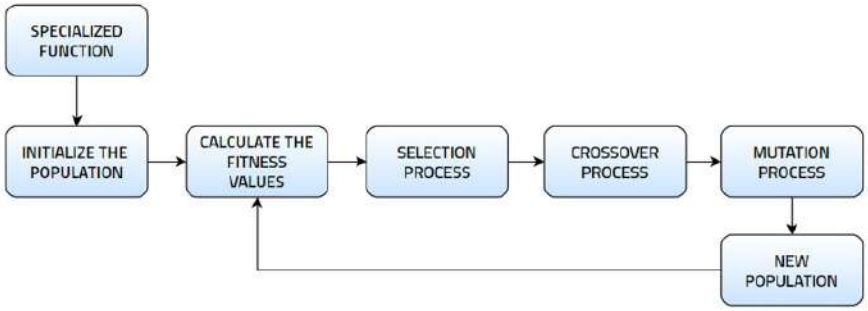

Figure 56: Key steps in a GA.

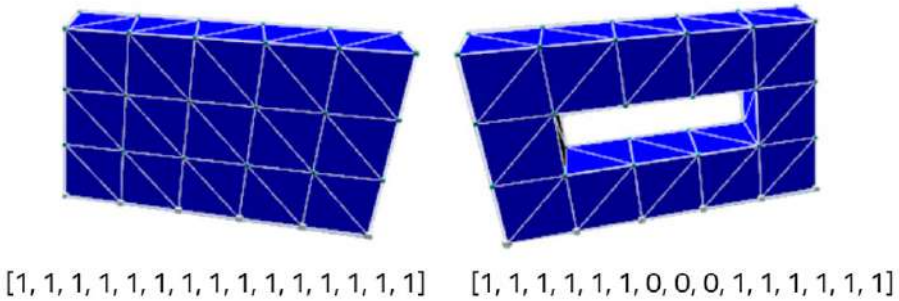

Figure 57: Encoded genes for individuals and their graphic representation.

To evaluate the behavior of GA-optimized structures, simulations were run using Project Chrono. The setup involves three steps: selecting the system (smooth or non-smooth contacts, both equivalent here due to lack of collisions), choosing a time-stepper

(Euler-Implicit-Linearized for its speed and FEA support), and picking a solver (Minres, for good convergence and FEA compatibility).

Fig. 58 illustrates all these steps, highlighting various options available for each component. The selected choices are shown in bold.

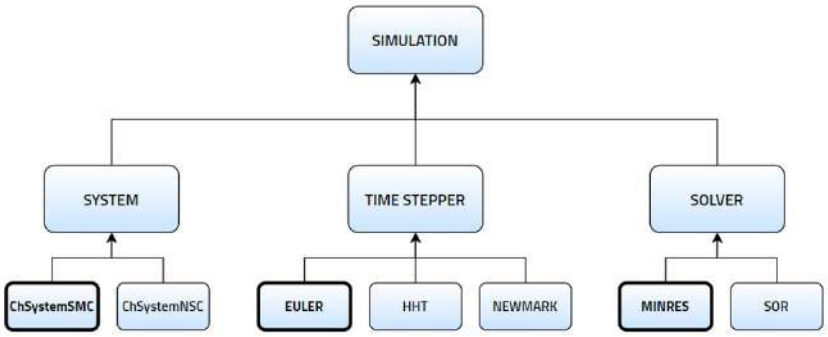

Figure 58: Steps for setting up a simulation.

The first approach used as an objective function a combination of the number of removed elements and the stress to which the structure was subjected:

$$F = 2^R * (max_S - max_{SS}) \tag{9}$$

While this led to some material minimization, as seen in Fig. 59, the overall structural optimization results were not satisfactory. This was primarily due to the exponential term, which excessively weighted the number of removed elements, causing the algorithm to oscillate between a limited number of individuals, thereby impeding its ability to explore for truly optimal designs.

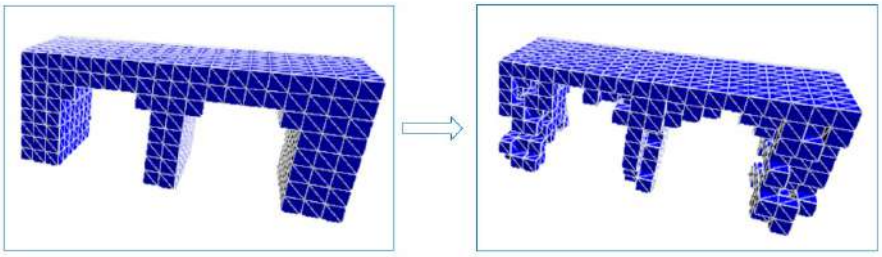

Figure 59: Resulted structure after the first tests.

Following the initial unsatisfactory results, it became evident that the fitness function required refinement. Tests were conducted using various fitness function formulations on a small, controlled structure ($10 \times 5 \times 3$ cubes along the O_x, O_y, and O_z axes, respectively) with an initial population of 100 individuals.

To simplify the function's notation, the definitions from Equation (10) are used.

$$S = max_S - max_{SS} \tag{10}$$

In Equation 10:

- max_S represents the maximum allowed stress.
- max_{SS} represents the maximum stress experienced by the structure.
- R represents the number of removed elements from the structure.

Table 13 presents the number of iterations required for the algorithm to achieve population stabilization, defined as the point where 10% of the individuals account for 90% of the selection probability for subsequent generations.

Table 13: The number of iterations for each fitness function

No.	Fitness function	Iterations
(1)	2^R*S	4
(2)	$R^{0.5}$*S	31
(3)	R*S	40
(4)	R^2*S	40

After testing and analyzing the fitness functions in Fig. 60, the last one was found to be the best. A key improvement observed in the resulting structures is the significant reduction of elements in the second layer by the algorithm, see some outputs in Fig. 61.

Figure 60: Evolution of fitness values for each objective function.

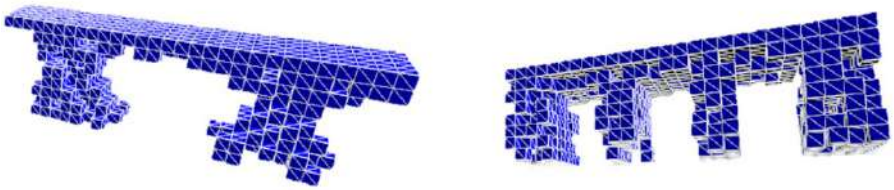

Figure 61: Structures following the last tests.

This approach successfully demonstrated the utility and robustness of GA for structural optimization. Our novel chromosome representation allows for efficient navigation

of large design spaces, providing notable computational benefits compared to conventional methods, especially for large-scale problems. This work highlights the practical integration of GAs as a powerful, automated solution, with potential for comparison against existing frameworks. The complete study [16] was presented at the C-category conference, namely the *IEEE International Symposium on Innovations in Intelligent Systems and Applications*.

II.2. Parametric optimization using genetic algorithms

Problem statement

In structural design, changes such as increased loads or new functional requirements often necessitate adapting existing designs. A common scenario involves a cantilever beam subjected to excessive transverse loading, leading to dangerous bending stresses. While over-dimensioning by increasing the cross-section's moment of inertia is a straightforward solution, it frequently results in material waste, particularly since the maximum bending moment is localized at the fixed end. A more efficient approach involves enhancing the beam's load-bearing capacity by adding a truss support without modifying the primary cross-section. This introduces a critical optimization problem: determining the optimal truss position to maximize load-bearing capacity up to a specified tension threshold, and comparing this solution against the conventional over-dimensioning method.

This section addresses this specific structural optimization challenge for a cantilever beam (2000 mm length, 40×20 mm rectangular section, generic steel, 1 kN end load) by focusing on determining the optimal horizontal and vertical distances for positioning the truss support ends.

There are three primary objectives of this research. The first one is to develop and implement a GA specifically tailored to solve this cantilever beam optimization problem. The second involves the validation of the GA's results through comparison with solutions obtained from other commercially available optimization tools. Lastly, the last objective is to evaluate the feasibility and material efficiency of the proposed GA-driven design solution relative to the straightforward approach of simply increasing the beam's cross-sectional moment of inertia.

Parametric optimization using genetic algorithm

The analytical formulation and detailed mathematical model for the structural optimization problem are comprehensively presented in the original publication [38]. Ultimately, the problem seeks to reduce the total mass of the system. This objective is quantified by the following function, which the genetic algorithm aims to minimize, see Equation (11).

$$m_{total} = \rho_{steel} \cdot h \cdot t \cdot (a + l_{23}) \tag{11}$$

In Equation (11):

- m_{total} represents the total mass of the system, in kilograms (kg).
- ρ_{steel} is the density of the generic steel material, given as 7859 kg/m^3.
- h is the section height, parallel to the load direction, with a value of 40 mm.
- t is the section width, perpendicular to the load direction, with a value of 20 mm.
- a is the length of the beam, specified as 2 m.
- l_{23} is the length of the truss.

Building upon this foundational understanding, the subsequent discussion in this section will primarily focus on the development and application of the genetic algorithm as the chosen methodology for solving this specific optimization challenge.

The GA's iterative process (visualized in Fig. 62) involves:

- Population Generation: Randomly initialized with 500 individuals in binary format, ensuring feasibility.
- Selection: Filters feasible individuals, computes fitness, Equation (11), applies elitism, and then uses one of several implemented mechanisms (Roulette, Truncation, Tournament, Boltzmann, Rank-based, Stochastic Universal Sampling) with specific parameters (e.g., Boltzmann's T and k, Tournament size $N_u = 30$, Truncation $k = 30$).
- Crossover: Employs uniform crossover, where offspring inherit genes from Parent 1 with a 50% chance of replacement by Parent 2's corresponding gene.
- Mutation: Applies bit-flip mutation with a probability that decreases with each generation $\left(p_{mut} = \frac{1}{g_c + 1}\right)$, balancing exploration and exploitation.
- Elitism: The best individual from the current generation replaces the least fit in the new population, preserving optimal solutions.
- Solution history: Superior solutions found are saved in a history list.
- The cycle is repeated throughout 100 generations.

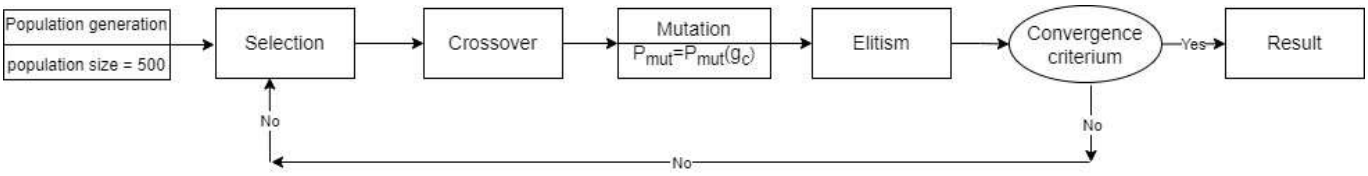

Figure 62: The flowchart of the genetic algorithm.

Genetic algorithms, by their stochastic nature, risk converging to a suboptimal solution rather than the true global optimum. To mitigate this, two primary strategies were explored:

1. *Config 1*: The algorithm is executed multiple times, and the best solution from these runs is either directly adopted or used to form the initial population for a final, refined execution.

2. *Config 2*: An initial run identifies a suboptimal region, after which the problem's search domain is narrowed to this vicinity. A subsequent algorithm run then explores this restricted area, significantly increasing the probability of converging to the global optimum. By design, this method involves running the algorithm at least twice in separate instances.

The algorithm's performance was rigorously evaluated across various configurations, with success measured by the frequency of convergence to the optimal solution and the average proximity of solutions within the design space (Error Rate, calculated via Root Mean Square Error, see Equation (12).

$$\text{Accuracy} = \frac{\text{Number of goods}}{\text{Number of iterations}} \times 100 \tag{12}$$

The evaluation setup involved:

1. All runs were consistently executed for 100 generations, with the mutation formula remaining constant.

2. A baseline configuration involved 10 consecutive runs for each of the six implemented selection methods (Roulette Wheel, Truncation, Tournament, Rank-based, Stochastic Universal Sampling, and Boltzmann), tested with population sizes of 500 and 1000 individuals.

3. The two additional algorithm configurations (as previously described) were evaluated with a population size of 100 individuals.

Analysis of the results (Tables 14 and 15) revealed that while larger populations generally tend to increase global optimum convergence, in this specific instance, doubling the population size from 500 to 1000 did not yield a significant improvement in accuracy or error function. This negligible gain in performance came at a substantial computational cost, increasing runtime by three to seven times.

The two algorithm configurations were evaluated using populations of 100 individuals each. For the domain narrowing approach, the solution from the initial run defined a refined search space, specifically a square with sides equal to 10% of the overall design space. This strategy inherently necessitates at least two algorithm executions. The results and observations from these evaluations are presented in Table 16 and Table 17.

The convergence of the GA to the global optimum is critically dependent on the chosen input configuration and selection method. Specifically, *Config 1*, when combined with Boltzmann, Truncation, Tournament, or Rank-based selection methods, consistently achieved a 100% success rate in converging to the global optimum, proving effective by

Table 14: Performance of selection methods on two population sizes: 500 and 1000 individuals

	Iteration	Boltzmann	Roulette	Truncation	Tournament	Rank based	SUS
500 individuals	1	17, 1200	18, 1201	18, 1200	17, 1200	17, 1200	28, 1200
	2	17, 1200	116, 1201	20, 1200	17, 1200	17, 1200	48, 1200
	3	17, 1200	32, 1201	17, 1200	17, 1200	17, 1200	17, 1200
	4	17, 1200	58, 1208	32, 1200	17, 1200	17, 1200	20, 1200
	5	17, 1200	38, 1200	17, 1200	17, 1200	17, 1200	20, 1200
	6	17, 1200	51, 1202	18, 1200	17, 1200	17, 1200	67, 1200
	7	17, 1200	30, 1204	30, 1200	17, 1200	17, 1200	32, 1200
	8	17, 1200	108, 1200	21, 1200	17, 1200	17, 1200	56, 1200
	9	17, 1200	70, 1200	21, 1200	17, 1200	17, 1200	20, 1200
	10	17, 1200	40, 1200	32, 1200	17, 1200	17, 1200	33, 1200
1000 individuals	1	17, 1200	36, 1201	72, 1201	17, 1200	17, 1200	17, 1200
	2	17, 1200	43, 1201	46, 1200	17, 1200	17, 1200	17, 1200
	3	17, 1200	86, 1200	19, 1207	17, 1200	17, 1200	25, 1200
	4	17, 1200	96, 1200	26, 1205	17, 1200	17, 1200	21, 1200
	5	17, 1200	73, 1200	58, 1200	17, 1200	17, 1200	23, 1200
	6	17, 1200	77, 1205	35, 1203	17, 1200	17, 1200	34, 1200
	7	17, 1200	21, 1200	21, 1204	17, 1200	17, 1200	35, 1200
	8	17, 1200	91, 1200	33, 1200	17, 1200	17, 1200	33, 1200
	9	17, 1200	40, 1201	22, 1203	17, 1200	17, 1200	24, 1200
	10	17, 1200	21, 1208	21, 1200	17, 1200	17, 1200	18, 1200

Table 15: The global optimum convergence for population of 500 and 1000 individuals

No individuals	Metric	Boltzmann	Roulette	Truncation	Tournament	Rank based	SUS
500	Accuracy	100%	0%	20%	100%	100%	20%
	RMSE	0	35,48	5,75	0	0	16,74
1000	Accuracy	100%	0%	0%	100%	100%	10%
	RMSE	0	35,28	17,85	0	0	7,05

Table 16: The results for both configurations

	Iteration	Boltzmann	Roulette	Truncation	Tournament	Rank based	SUS
Config 1	1	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	65, 1200
	2	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	24, 1200
	3	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	17, 1200
	4	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	17, 1200
	5	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	17, 1200
	6	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	17, 1200
	7	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	17, 1200
	8	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	17, 1200
	9	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	17, 1200
	10	17, 1200	25, 1201	17, 1200	17, 1200	17, 1200	17, 1200
Config 2	1	17, 1200	18, 1200	18, 1200	20, 1200	32, 1200	17, 1200
	2	72, 1200	18, 1200	32, 1200	22, 1200	18, 1200	17, 1200
	3	20, 1200	24, 1200	24, 1200	18, 1200	20, 1200	18, 1200
	4	32, 1200	32, 1200	20, 1200	18, 1200	27, 1200	32, 1200
	5	24, 1200	17, 1200	32, 1200	24, 1200	17, 1200	20, 1200
	6	34, 1200	17, 1200	18, 1200	18, 1200	18, 1200	33, 1200
	7	17, 1200	18, 1200	32, 1200	17, 1200	17, 1200	17, 1200
	8	20, 1200	17, 1200	24, 1201	32, 1200	20, 1200	32, 1200
	9	17, 1200	20, 1200	17, 1200	32, 1200	33, 1200	17, 1200
	10	18, 1200	65, 1200	20, 1200	32, 1200	17, 1200	65, 1201

Table 17: The global optimum convergence for both configurations

Config	Metric	Boltzmann	Roulette	Truncation	Tournament	Rank based	SUS
Config 1	Accuracy	100%	0%	100%	100%	100%	80%
	RMSE	0	5,70	0	0	0	10,84
Config 2	Accuracy	30%	30%	10%	10%	30%	40%
	RMSE	13,42	11,37	6,3	6,16	5,48	12,29

maintaining population diversity and balancing exploration with exploitation. Conversely, the roulette wheel selection method frequently failed to converge to the global minimum, primarily due to a rapid loss of population diversity and an imbalance between explo-

ration and exploitation. While the optimized truss system was marginally heavier (0.24 kg) than the over-dimensioned beam at the initial tested load, further analysis revealed its significant advantage at higher loads: as the load increases, the weight of the baseline over-dimensioned model escalates disproportionately more than that of the optimized truss concept, highlighting the long-term efficiency and scalability benefits of the optimized design. Regarding *Config 2*, its relatively poor performance, despite the theoretical advantage of narrowing the feasible design space, suggests that a higher degree of exploratory behavior is necessary to fully exploit the identified suboptimal subspace. While this configuration offers the main advantage of requiring a reduced number of runs, strictly evaluating the solution's real-world feasibility will necessitate addressing current simplifications.

This work, in its original form, has been presented at a C-ranking conference, specifically the *Proceedings of the International Conferences on Applied Computing* [38].

III. AI for medical triage and clinical decision support

Within this chapter dedicated to AI for monitoring, diagnosis, and optimization in complex and interdisciplinary systems, diverse applications of artificial intelligence in critical domains are explored. An essential component of this exploration is AI for medical triage and clinical decision support, where AI's potential to revolutionize medical care is huge. The study [172], co-authored by the author of this thesis, highlights contributions in this field. This research specifically investigates the utilization of machine learning for medical triage, and its primary conclusions, demonstrating the efficacy of the proposed methodologies, will be presented in detail.

The main elements identified and supported by the study behind this section include:

- A. Vântu, A. Vasilescu, **A. Băicoianu**, 2023 - *Medical emergency department triage data processing using a machine-learning solution* [172]

  Main findings:

  - The NN-Sequential algorithm performed better in the simplified problem.

    The initial problem involved classifying patients into 5 levels of the Emergency Severity Index (ESI), which presented challenges due to imbalanced data distributions across these levels. To address this, the classification task was simplified to a 3-class problem: critical, urgent, and non-urgent. In this simplified scenario, the NN-Sequential algorithm demonstrated superior performance compared to the other models, indicating that reducing the complexity of the classification target, especially when dealing with data imbalance, can significantly improve model efficacy.

  - Overall, in all experiments, the NN-Sequential model yielded superior results.

    Across all the experiments conducted, which included both the original 5-class ESI classification and the simplified 3-class classification, the NN-Sequential

model consistently outperformed the Logistic Regression and Random Forest Tree models. In the initial experiment (likely the 5-class ESI), the NN-Sequential model achieved ROC-AUC scores for each ESI emergency code of 0.59%, 0.76%, 0.71%, 0.78%, and 0.64%. This shows varying performance across different ESI levels, with some struggling more than others (e.g., the 0.59% for potentially critical Level 1 patients). After applying methods to balance the data, the model's ROC-AUC scores for each emergency code improved to 0.72%, 0.75%, 0.69%, 0.74%, and 0.78%. This highlights the positive impact of data balancing on the model's generalization capabilities. Furthermore, in the final experiment, which involved the three-class classification problem, the NN-Sequential algorithm achieved strong ROC-AUC scores of 0.76%, 0.72%, and 0.84% for the respective emergency codes. While the Random Forest Tree model showed comparable metric outcomes in this specific simplified problem, the NN-Sequential model generally maintained its lead. Despite its overall strong performance, the model still faced challenges in accurately classifying the most critical Level 1 patients, as indicated by the lower initial ROC-AUC score for that level.

– As part of the study's practical application, a web-based application was developed. This tool was designed to improve the triage process by providing an interface for patient overview and early monitoring. Clinicians found the tool easy to use and interacted with it reliably, indicating that the technology is user-friendly and suitable for integration into clinical workflows, supporting its practical implementation.

Problem statement and dataset

Machine learning offers diverse algorithms with practical solutions increasingly impacting quality of life. Artificial intelligence applications span from renewable energy systems [154] to critical improvements in Emergency Departments (EDs) [167], [149]. Recent decades have seen a surge in research on ML in medicine, with studies from 2001 to 2022 [155] revealing specific medical data processing methods. A key challenge lies in efficiently accumulating and recalling medical records to accurately evaluate and classify symptoms in emergency medicine, a process significantly enhanced by ML algorithms.

The Emergency Department represents one of the most critical decision points in any healthcare system, primarily operating through a triage algorithm designed to prioritize emergencies. While triage protocols vary globally, in Romania, the system is based on the Emergency Severity Index (ESI) algorithm, initially developed in 1999 [65]. This protocol categorizes patients into five distinct emergency codes:

1. Level 1 (Red): The most critical level, demanding immediate intervention due to a life-or-death situation.

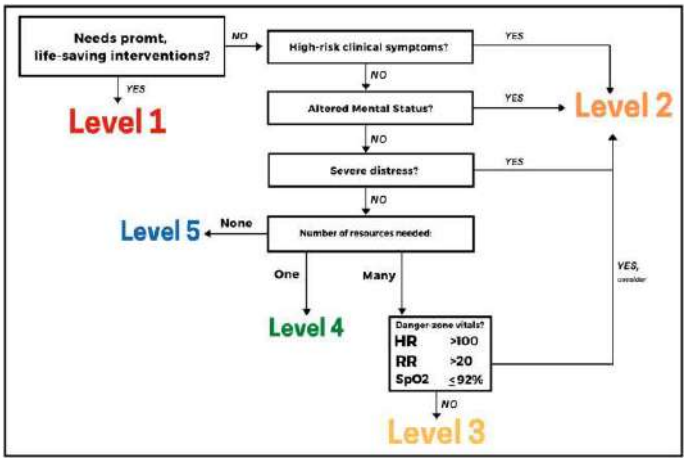

Figure 63: ESI triage protocol. [99]

2. Level 2 (Yellow): Indicates a high risk of deterioration, allowing for a wait time of up to 10 minutes.

3. Level 3 (Green): Requires urgent care, with patients able to wait up to 30 minutes without immediate life-threatening risks.

4. Level 4 (Blue): Represents a stable health status requiring only one resource, considered non-urgent, with a potential wait of up to one hour.

5. Level 5 (White): Also non-urgent, where patients may wait for up to two hours.

Each patient follows the triage protocol represented here in Fig. 63, and a trained clinician assigns an emergency code based on specific rules, including chief complaints, age, past medical history, vital signs, and current treatment, following a series of triage questions.

However, a pervasive issue in Emergency Departments is overcrowding. Although the ESI triage protocol aims for accuracy, it remains susceptible to human error, particularly when medical staff are overwhelmed by a large influx of patients. This susceptibility can lead to either overestimation or underestimation of a patient's health status. While overestimation may result in an inefficient allocation of medical resources, underestimation carries severe consequences, potentially leading to negative outcomes, including increased mortality rates [190]. This section addresses the multi-class classification problem of predicting emergency codes for individuals within the specific context of the Romanian emergency medical system, aiming to mitigate these challenges through a machine learning approach.

The dataset employed in this section [74] consists of adult visits to the Emergency Department (ED), gathered from three different facilities between March 2014 and July 2017. Although originally intended for predicting patient admission outcomes [75], the dataset's features are highly applicable to emergency level classification tasks.

The raw dataset contained 972 features and 560486 entries, stored in a 1.32GB CSV file. These features were broadly categorized into demographics, past medical history,

Table 18: Selected dataset feature categories

Category	Attribute	Data type
Response variable	esi	Categorical
Past medical history	From 2sndarymalig to whtblooddx	Categorical
Demographics	age gender arrivalmode previousdispo	Numeric Categorical Categorical Categorical
Triage Vital Signs	triage_vital_hr triage_vital_rr triage_vital_sbp triage_vital_dbp triage_vital_o2 triage_vital_temp	Numeric Numeric Numeric Numeric Numeric Numeric
Chief Complaints	From cc_abdominalcramping to cc_wristpain	Categorical

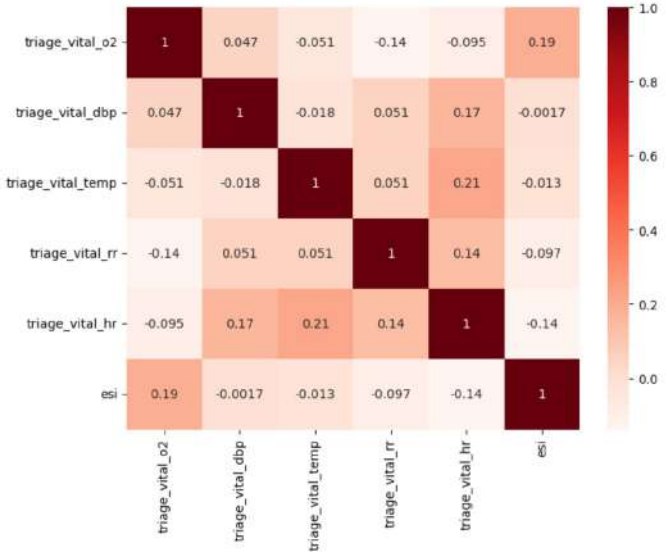

Figure 64: ESI - Vital Triage Signs correlation matrix.

historical vitals, triage vital signs, outpatient medication, chief complaints, historical labs, imaging, and hospital usage.

For this study, a focused subset of 409 features was selected, retaining all original entries. This selection prioritized attributes most suitable for the Romanian ED patient file, including demographics, past medical history, triage vital signs, and chief complaints. The resulting dataset is a mix of numerical, binary, and text values. Data types were determined for efficient workflow, as summarized in Table 18.

Initial exploratory data analysis involved understanding variables, assessing them, and plotting correlation matrices to identify relationships between features and the *esi* response variable. Correlations with triage vital signs, past medical history, and chief complaints were specifically analyzed, as depicted in Fig. 64, Fig. 65, and Fig. 66 respectively. For past medical history and chief complaints, ten most relevant features were selected for each category.

The dataset was then split into training (80%) and test (20%) sets using the *sklearn* library in Python. Both sets underwent a comprehensive preprocessing pipeline, including

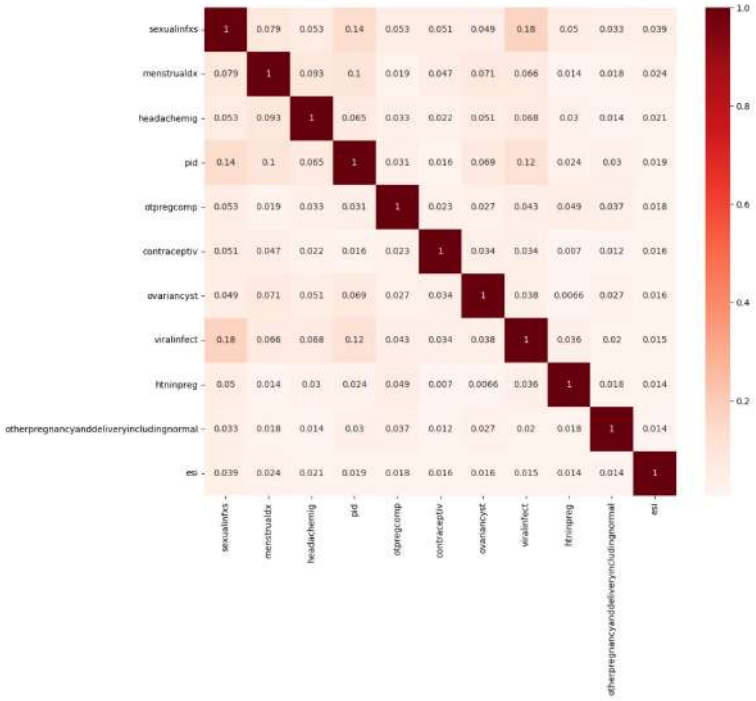

Figure 65: ESI - Past Medical History correlation matrix.

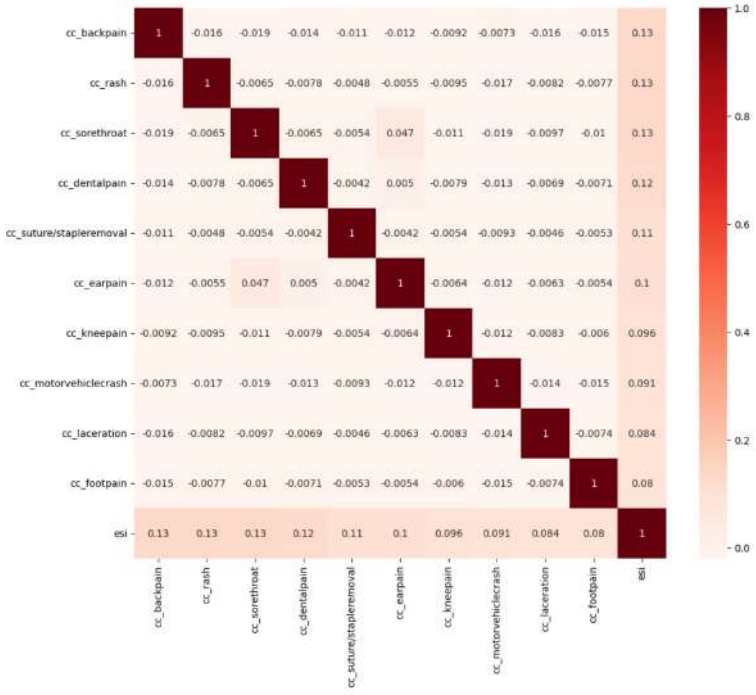

Figure 66: ESI - Chief Complaints correlation matrix.

handling missing values, data scaling, and encoding text attributes.

Missing values were identified in features such as *esi*, *age*, *arrivalmode*, and various *vital_triage* signs, along with some chief complaints. Entries with missing *esi* values were dropped, as imputing this categorical response variable could misrepresent ground truth. Records missing *esi* also had missing values in other listed columns, thus resolving those as well. Missing *age* values were similarly treated by dropping the few records, as it would not significantly affect the model's prediction outcome. For continuous numerical attributes like triage vital signs, missing values were imputed with the median using *SimpleImputer* from *sklearn.impute* package. For *arrivalmode*, the most frequent strategy was applied for imputation.

All numerical attributes were scaled using *StandardScaler()* from *sklearn*. This step was crucial for model performance, as the dataset contained 7 numerical features with values starting from 18 and going up to 105, while other features were binary categorical (0 or 1). Meanwhile, categorical text features, specifically *gender*, *arrivalmode*, and *previousdispo*, were transformed into numerical format using *OneHotEncoder*. This process converts each unique categorical value into a new binary feature, enhancing compatibility with machine learning models.

Methodology and experimental evaluation

This research addresses a multi-class classification problem: predicting a patient's emergency code. To achieve this, three supervised machine learning models were selected for evaluation:

- *Neural Network Models:* Specifically, the NN Sequential Model, chosen for its capability in exploring complex data.
- *Decision Trees:* Represented by the Random Forest Tree Algorithm, which offers a robust and comparatively "cheap" approach with good predictive capabilities due to its ensemble nature.
- *Regression Algorithms:* The Logistic Regression Algorithm was chosen as it is a regression technique adaptable for classification problems, known for being a good initial estimator and widely used in medical data processing.

For a comprehensive evaluation of model performance, especially critical in medical contexts, a suite of metrics beyond mere accuracy was employed. These include the ROC (Receiver Operating Characteristic)-AUC (Area Under the ROC curve) score, precision, recall, specificity, and the confusion matrix. These metrics provide deeper insights into the model's behavior and are particularly relevant for clinical decision-making.

The computational experiments evaluated the performance of the chosen ML models (NN-Sequential, Random Forest Tree, and Logistic Regression) across different data handling strategies. These strategies included using the original imbalanced data, and two approaches to handle imbalance: *SMOTE* and *ADASYN* oversampling techniques.

Table 19: Models hyper-parameters

Model	Hyper-parameter	Value
Logistic Regression	multi_class	multinomial
	solver	lbfgs
	C	10
NN-Sequential	loss	categorical_crossentropy
	optimizer	SGD
	metrics	accuracy
	epochs	100
	batch_size	128
	callback	EarlyStopping
Random Fores Tree	n_estimators	50

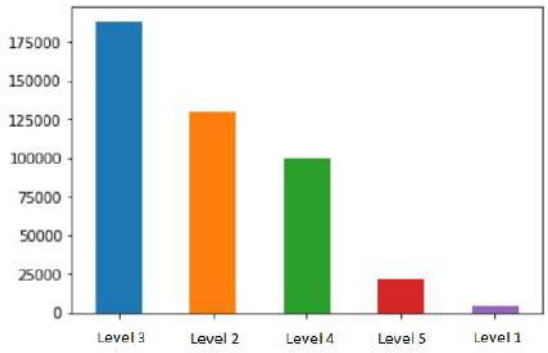

Figure 67: The distribution of target classes.

Model hyperparameters were empirically tuned through a trial-and-error process to optimize performance. For Logistic Regression, a `multinomial` option with `lbfgs` solver and C=10 was used. The NN-Sequential model employed `categorical_crossentropy` as the loss function, an SGD optimizer (with specific learning rate, decay, momentum, and nesterov parameters), trained for 100 epochs with a batch size of 128 and an EarlyStopping callback. Its architecture consisted of 4 hidden layers with `relu` activation and 30 nodes, and an output layer with `softmax` activation and 5 nodes. The Random Forest Tree model utilized 50 estimators. These final hyperparameters are detailed in Table 19.

Computational efficiency, including training and testing time complexity, and memory usage, was considered for each model [172]. The Logistic Regression model exhibits linear complexity with respect to data size and features. The Random Forest Tree, while having a more complex training time, is generally comparatively faster. The NN-Sequential model's memory footprint is determined by its total number of parameters (e.g., 12780 parameters for the initial layers in this study), with its time complexity being more challenging to precisely quantify from standard library implementations.

The dataset exhibits a significant class imbalance, as illustrated in Fig. 67. The majority class is *Level 3*, with 188,245 entries, while *Level 1* is a minority class with only 4,265 entries. This imbalance poses a challenge, as ML models are prone to performing poorly on minority classes [91].

Experiments were conducted using all three models, initially on the raw, imbalanced dataset (80% train-20% test split).

Some details about the results on *imbalanced data (raw data)*:

Figure 68: Confusion matrix of the NN Sequential model.

Table 20: Per-Class performance metrics of the NN Sequential model - Raw data

Classes	Recall	Precision	Specificity	ROC-AUC
Level 1	0.18	0.71	0.99	0.59
Level 2	0.65	0.69	0.92	0.76
Level 3	0.71	0.65	0.72	0.71
Level 4	0.67	0.65	0.89	0.78
Level 5	0.29	0.54	0.97	0.64

- *NN Sequential model (Experiment 1):*
  - Trained for 41 epochs (out of 100 max) before early stopping, achieving an accuracy of 0.67.
  - The confusion matrix is shown in Fig. 68, and detailed metrics per class are in Table 20.
  - *Key findings: Level 1* showed low recall (0.18), indicating many patients needing *Level 1* care were misclassified. However, its precision (0.71) and specificity (0.99) were high. *Levels 2, 3, and 4* demonstrated the best overall precision and recall scores. The model struggled to distinguish between *Level 2* and *Level 3*, as seen in the confusion matrix. *Level 4* had the highest ROC-AUC (0.78).
- *Random Forest Tree (Experiment 2):*
  - Utilized 50 estimators, determined via *GridSearchCV*. The confusion matrix is presented in Fig. 69.
  - *Key findings:* Metrics in Table 21 show this model performed generally less well than the NN Sequential model, with lower precision and recall across most classes. *Level 1* recall remained 0.18, but precision dropped to 0.42.

Table 21: Per-Class performance metrics of RFT - Raw data

Classes	Recall	Precision	Specificity	ROC-AUC
Level 1	0.18	0.42	0.99	0.59
Level 2	0.61	0.63	0.85	0.73
Level 3	0.68	0.62	0.69	0.69
Level 4	0.61	0.62	0.89	0.75
Level 5	0.23	0.48	0.98	0.61

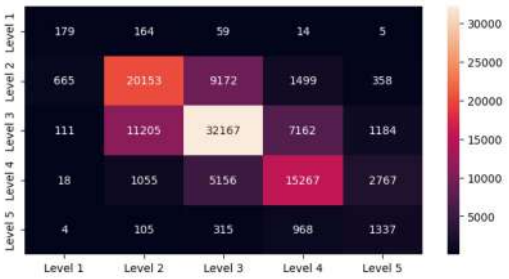

Figure 69: Confusion matrix of RFT.

Figure 70: Confusion matrix of LR.

- *Logistic Regression (Experiment 3):*
  - Configured with `multi_class='multinomial'` and `solver='lbfgs'`. The confusion matrix is in Fig. 70.
  - *Key findings:* Table 22 shows its performance. *Level 4* had the lowest recall for this model. Its ROC-AUC was similar to Random Forest Tree, and precision was slightly lower than NN Sequential. Overall, its average performance was slightly behind the NN model, except for specificity (average 0.9).

Among the three classifiers evaluated on the imbalanced dataset, the NN Sequential model achieved the most favorable overall performance. However, a persistent issue across all models was the notably low recall for *Level 1* (0.18), reflecting a significant failure to correctly identify many patients in the most critical emergency category. This outcome underscores the challenges posed by class imbalance, particularly for minority class prediction, where precision and recall are essential performance indicators.

As previously noted, the dataset exhibits a significant class imbalance, with *Level 3* being the majority (42%) and *Level 1* the minority (1%), as detailed in Fig. 67. To address this challenge in multi-class classification, *SMOTE* (Synthetic Minority Over-sampling Technique) [57] was employed. Specifically, *SMOTE*-NC (for Nominal and Continuous

Table 22: Per-Class performance metrics of LR - Raw data

Classes	Recall	Precision	Specificity	ROC-AUC
Level 1	0.18	0.68	0.99	0.59
Level 2	0.64	0.67	0.87	0.75
Level 3	0.68	0.64	0.78	0.70
Level 4	0.69	0.62	0.89	0.78
Level 5	0.14	0.55	0.99	0.56

Table 23: Per-Class performance metrics of LR - SMOTE balanced data

Classes	Recall	Precision	Specificity	ROC-AUC
Level 1	0.42	0.08	0.95	0.69
Level 2	0.55	0.65	0.87	0.71
Level 3	0.59	0.67	0.79	0.69
Level 4	0.64	0.59	0.87	0.75
Level 5	0.42	0.24	0.97	0.67

Table 24: Per-Class performance metrics of the NN Sequential - SMOTE balanced data

Classes	Recall	Precision	Specificity	ROC-AUC
Level 1	0.35	0.14	0.98	0.66
Level 2	0.65	0.66	0.96	0.75
Level 3	0.63	0.68	0.78	0.70
Level 4	0.67	0.62	0.89	0.78
Level 5	0.39	0.33	0.95	0.67

features) was used to handle the mixed data types, generating synthetic samples for minority classes. This process balanced the dataset, increasing the number of entries for each level to approximately 188,242.

All three models were re-evaluated on this *SMOTE*-balanced dataset, keeping their hyperparameters constant.

Some details about the results on *SMOTE-balanced dataset*:

- Logistic Regression with *SMOTE*: Showed increased *recall* for minority classes (*Level 1* and *Level 5*), but *Level 1*'s *precision* significantly dropped (0.08), indicating a high rate of false positives for critical cases (Table 23).
- NN Sequential Model with *SMOTE*: Achieved an improved average ROC-AUC (0.71) compared to its performance on imbalanced data. *Level 1 recall* improved, but *precision* remained low (0.14) (Table 24).
- Random Forest Tree with *SMOTE*: Generally showed lower *recall* scores across *Levels 1, 3, 4, and 5* compared to the other models using *SMOTE*, and its overall ROC-AUC scores were also lower than its performance on the original imbalanced data (Table 25).

Despite *SMOTE*'s ability to balance the dataset and increase data volume, the improvements in model performance were not consistently significant across all metrics or models. This suggests that while synthetic sampling can address class imbalance, it does not guarantee substantial gains in prediction accuracy, especially for minority classes where *precision* remained a persistent challenge.

The original study also explored the use of *ADASYN* for handling imbalanced data. For a detailed discussion, refer to [172]. Additionally, a simplified classification problem

Table 25: Per-Class performance metrics of RFT - SMOTE balanced data

Classes	Recall	Precision	Specificity	ROC-AUC
Level 1	0.24	0.16	0.98	0.61
Level 2	0.60	0.58	0.74	0.71
Level 3	0.57	0.63	0.75	0.66
Level 4	0.61	0.54	0.85	0.73
Level 5	0.29	0.31	0.96	0.63

was investigated in an effort to improve prediction performance. Instead of the original 5-level classification, the problem was reformulated into a 3-level classification. This involved grouping *Level 1* and *Level 2* into a `Critical` class, *Level 3* remaining as `Urgent`, and *Level 4* and *Level 5* being combined into a `Non-urgent` class. This grouping reflects a natural prioritization, where levels above *Level 3* represent critical situations requiring maximum priority, and those below signify non-urgent cases. The resulting class distribution for this simplified problem was: `Critical` - 30%, `Urgent` - 43%, and `Non-urgent` - 27%.

For the simplified 3-level classification problem, the three models were re-evaluated. The NN-Sequential model consistently performed best, particularly in terms of *recall* across the `Critical`, `Urgent`, and `Non-urgent` classes, see details in [172]. The Random Forest Tree model showed good performance but was slightly inferior to the NN-Sequential. In contrast, the Logistic Regression model performed poorly. This simplified approach generally yielded better results for the NN-Sequential model. The study also compared its findings (`eUPU`) with other relevant research in medical triage and admission prediction (KTAS [39], KATE [78], ED-Admission [74]), highlighting the potential for further improvements with more advanced models like XGBoost [172].

This study offers key conclusions on the application of ML in emergency medical triage, aligning with the broader objectives of AI for monitoring, diagnosis, and optimization in complex and interdisciplinary systems. Specifically, it explored a ML approach to automate emergency code assignment after triage, demonstrating the potential of supervised models to improve early medical diagnosis and streamline record processing. The study highlighted how AI can significantly enhance the efficiency of emergency triage, justifying real-time application and dynamic evaluation of these models. Significant challenges were encountered, including the difficulty of obtaining and customizing medical data for the Romanian context, handling missing values, and addressing severe class imbalance. To mitigate the impact of imbalanced data, the problem was simplified to a 3-level classification, where the NN-Sequential algorithm consistently demonstrated the best performance.

While this study offers valuable insights, significant opportunities for further improvement exist. It is important to note that the original research was published in *Heliyon*, a Q2-ranked journal, and has garnered numerous citations from both computer science and medical fields, underscoring the practical significance of its findings. Future work will involve collecting more comprehensive data and, importantly, integrating expert feedback. This feedback loop is facilitated by the developed web application, whose details, including a repository, are available in the original publication [172]. This tool serves as a user-friendly interface for patient overview and early monitoring. Initial findings indicate that clinicians found the tool reliable. Ultimately, achieving highly accurate prediction of emergency codes through such integrated systems will pave the way for creating truly digital and smart patient files.

IV. AI-driven accuracy optimization in environmental time series forecasting

In the broader context of using AI for improving the performance and reliability of monitoring systems, the following contribution explores how time series prediction accuracy can be enhanced through novel data augmentation techniques [17]. By employing fractal interpolation strategies in conjunction with LSTM models, the study targets a crucial challenge in the field: generating high-quality synthetic data that supports more accurate and robust predictions. This approach illustrates a practical blend of machine learning, optimization, and signal enhancement, offering valuable insights for applications in remote sensing and environmental monitoring.

The key contributions proposed and demonstrated by the study supporting this section are as follows:

- **A. Băicoianu**, C.G. Gavrilă, C. Păcurar, V.D. Păcurar, 2024 - *Fractal interpolation in the context of prediction accuracy optimization* [17]

  Main findings:

  - The study introduces and defines three distinct strategies for fractal-based time series augmentation:
    - *Closest Hurst Strategy* (CHS): maintains a constant Hurst exponent to preserve the complexity of the original signal.
    - *Closest Values Strategy (CVS)*: ensures minimal variation in interpolation while avoiding linearity.
    - *Formula Strategy (FS)*: uses a geometric formula to define vertical scaling, eliminating the need for iterative optimization.

    All strategies are applied during the preprocessing stage of LSTM model training, by generating n intermediate points between original time series values to increase dataset density.
  - The proposed methods were validated on four public datasets and one private meteorological dataset from Brașov, Romania, with a sampling frequency of 10 minutes, allowing for a fine-grained analysis of interpolation effectiveness.
  - Results show that LSTM models trained on fractally interpolated datasets achieved higher prediction accuracy than those trained on the original datasets with the same dimensionality.

    For the CHS, which is based on maintaining a constant Hurst exponent, the study successfully reduced the number of iterations required to stabilize the vertical scaling factor from 200 to just 15 by employing an optimization framework (e.g., `Optuna`). This represents a substantial efficiency gain in applying this strategy. The CVS was developed to address limitations of CHS, offering a non-linear interpolation with small variations. The FS, despite potentially introducing the highest level of noise, provides a significant advantage by not

requiring per-step optimization, as its vertical scaling factor is determined geometrically.

– The findings suggest a promising alternative to increasing sensor sampling rates, enabling both storage efficiency and improved model accuracy in resource-constrained monitoring systems.

Problem statement

The development of AI models requires access to large volumes of high-quality data. The performance of supervised learning algorithms depends directly on both the quantity and diversity of the data used during the training phase. However, obtaining such datasets can be costly or even impossible, especially in the case of historical time series data.

Time series represent a category of data characterized by sequential observations that describe a phenomenon or process where temporal order is essential. Classical interpolation methods are often employed in the preprocessing stage of predictive modeling. However, unlike domains such as image processing, there is no widely accepted standard for data augmentation in time series forecasting. Fractal interpolation has found diverse applications, from computer graphics [107] and image compression [111] to signal processing [121] and epidemic curve reconstruction. Its combination with machine learning, particularly in areas like stock trend prediction with SVM [177] or short-term prediction with Time-Delayed Neural Networks [188], underscores its growing relevance.

The broader field of synthetic data generation and time-series data augmentation has seen significant research interest, with comprehensive reviews such as [183] highlighting the importance of these techniques. Interpolation, in general, has proven to be a robust method for time series augmentation [125] and for reconstructing incomplete datasets [35]. While classical interpolation methods are widely used in ML prediction [24], this study specifically uses the unique properties of fractal interpolation to better represent real-world complexity and improve forecasting accuracy [22], [142].

In [142], a comparative analysis is conducted on the impact of fractal interpolation on the prediction accuracy of Long Short-Term Memory (LSTM) models, contrasted with traditional linear interpolation. Fractal interpolation generates intermediate data points within a set $\Delta = \{(x_i, y_i) \mid i \in \{0, 1, \ldots, n\}\}$, with its defining characteristic being the jagged, non-smooth output. While it preserves continuity, the resulting function is non-differentiable at any point. This property makes fractal interpolation particularly suitable for generating realistic synthetic data in domains where noise is intrinsic, such as sensor readings, audio signals, or EEG measurements.

The current study builds on the research initiated in [142] by addressing several open development directions. Specifically, it proposes three distinct fractal interpolation strategies for data preprocessing: the *Closest Hurst Strategy*, the *Closest Values Strategy*, and the *Formula Strategy*. While CHS builds upon approaches similar to those presented in [142], this work significantly advances it by addressing the challenge of optimally selecting

vertical scaling factors and demonstrating an optimized prediction model using the `Optuna` framework. Importantly, CVS and FS are novel approaches not previously explored in the literature.

The utility of these strategies extends beyond theoretical advancements. The methods are validated using both public datasets and a real-world meteorological dataset from Braşov, Romania. This approach demonstrates the potential to generate high-resolution data from lower-quality measurements, providing a solution for scenarios where sensor data is sparse, unreliable, or requires improved time resolution.

Methodology and experimental evaluation

This section outlines the methodology originality employed in [17], detailing the crucial steps that lead to the final findings, as visually represented in Fig. 71. The core contribution of this study is concentrated on the interpolation step within a time series prediction pipeline. Our primary objective is to evaluate the impact of data augmentation via fractal interpolation on the quality and accuracy of predictions. For this purpose, an LSTM model was selected due to its demonstrated efficacy in handling long-term dependencies, seasonal patterns, and missing data prevalent in climate time series, as well as its adaptability to dynamic environmental conditions. Each block illustrated in the diagram will be further explored in subsequent sections to highlight its specific role in the overall process.

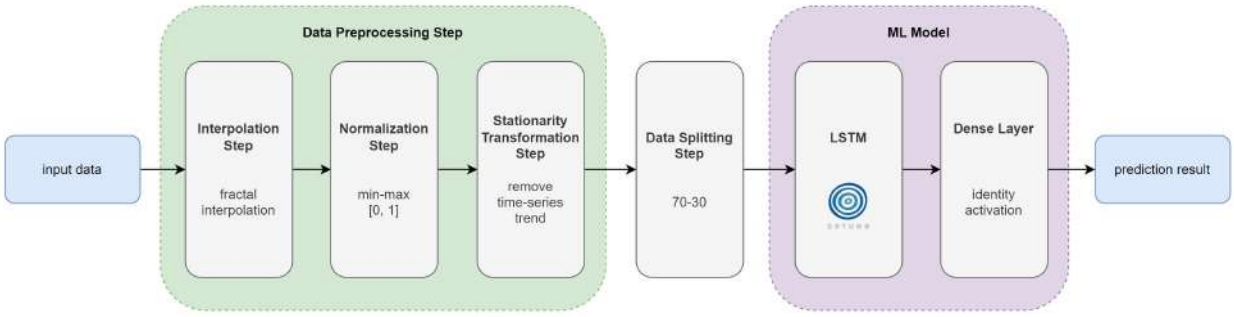

Figure 71: Fractal interpolation and LSTM forecasting methodology.

Data preprocessing step

This study utilizes two main categories of datasets for experimental validation:

- Meteorological data: A high-granularity dataset comprising 13,105 temperature entries recorded every 10 minutes from September 1 to November 30, 2021. This data, provided by the Forest Meteorology-Climatology Laboratory, *Transilvania University* of Brasov, was collected from a HOBO®RX3000 automatic weather station near Brașov, Romania. Formatted as a .csv file (385 KB), this dataset's fine temporal resolution (10 minutes) is really important for evaluating the ability of the proposed fractal interpolation strategies to reproduce relevant intermediate values and enhance data time resolution for predictive modeling.

- Additional public datasets: For comparative analysis and to align with previous research [142], four public datasets were also employed: Shampoo sales (36 data points, [81]), Airline passengers (144 data points, [81]), Annual wheat yields in Austria (57 data points, [55]), Annual maize yields in Austria (58 data points, [55]).

A time series is given as a set of data points as $\{(u_m, v_m) \in \mathbb{R}^2\} : m = 0, 1, \ldots, M$. The interpolation is then applied to a subset of those data points, i.e., interpolation points $\{(x_i, y_i) \in \mathbb{R}^2 : i = 0, 1, \ldots, N\}$. Both sets are linearly ordered with respect to their abscissa, meaning $u_0 < u_1 < \cdots < u_M$ and $x_0 < x_1 < \cdots < x_N$.

The interpolation step is an important component of this methodology, serving as a primary method for augmenting time series data by generating new data points within existing trends.

Let the family of functions $f_n : [x_0, x_N] \times \mathbb{R} \to [x_0, x_N] \times \mathbb{R}$ defined as

$$f_n \begin{bmatrix} x \\ y \end{bmatrix} = \begin{bmatrix} a_n & 0 \\ d_n & s_n \end{bmatrix} \begin{bmatrix} x \\ y \end{bmatrix} + \begin{bmatrix} c_n \\ e_n \end{bmatrix}, \tag{13}$$

which is constrained to

$$f_n \begin{bmatrix} x_n \\ y_n \end{bmatrix} = \begin{bmatrix} x_{n-1} \\ y_{n-1} \end{bmatrix} \quad \text{and} \quad f_n \begin{bmatrix} x_N \\ y_N \end{bmatrix} = \begin{bmatrix} x_n \\ y_n \end{bmatrix}$$

for every $n \in \{1, 2, \ldots, N\}$.

Consistent with the methodology presented in [108], solving these equations yields:

$$\begin{aligned} a_n &= \frac{x_n - x_{n-1}}{x_N - x_0}, \\ c_n &= \frac{x_N x_{n-1} - x_0 x_n}{x_N - x_0}, \\ d_n &= \frac{y_n - y_{n-1}}{x_N - x_0} - s_n \frac{y_N - y_0}{x_N - x_0}, \\ e_n &= \frac{x_N y_{n-1} - x_0 y_n}{x_N - x_0} - s_n \frac{x_N y_0 - x_0 y_N}{x_N - x_0}. \end{aligned} \tag{14}$$

The real coefficients a_n, c_n, d_n, e_n are established by the interpolation points. Additionally, s_n serves as a free vertical scaling factor, bounded by the condition $|s_n| < 1$.

This process involves two main substeps. First substep, the time series is divided into overlapping `sequence_size` subsets, which are later reunited. The second substep, the three distinct fractal interpolation strategies are introduced: the *Closest Hurst Strategy*, the *Closest Values Strategy*, and the *Formula Strategy*. These strategies aim to validate previous research and enhance interpolation techniques. Their application is demonstrated across both public datasets (Shampoo, Airline, Wheat, Maize) and the specific hourly meteorological dataset from Braşov.

Fig. 72 illustrates the general interpolation flow, while Algorithm 3 provides the pseudocode for the fractal interpolation computation. The description of the parameters for the

FRACTAL_INTERPOLATION procedure are:

- `subset`: subset created from original data as described in the first subset.
- s_n: a vector of vertical scaling factors which dictate how jagged (and fractal) will be the aspect of the generated data, in the sense that, as its name states, it scales the points vertically.
- `n_interpolation`: the number of distinct interpolation points to be generated between every 2 points of the original data.

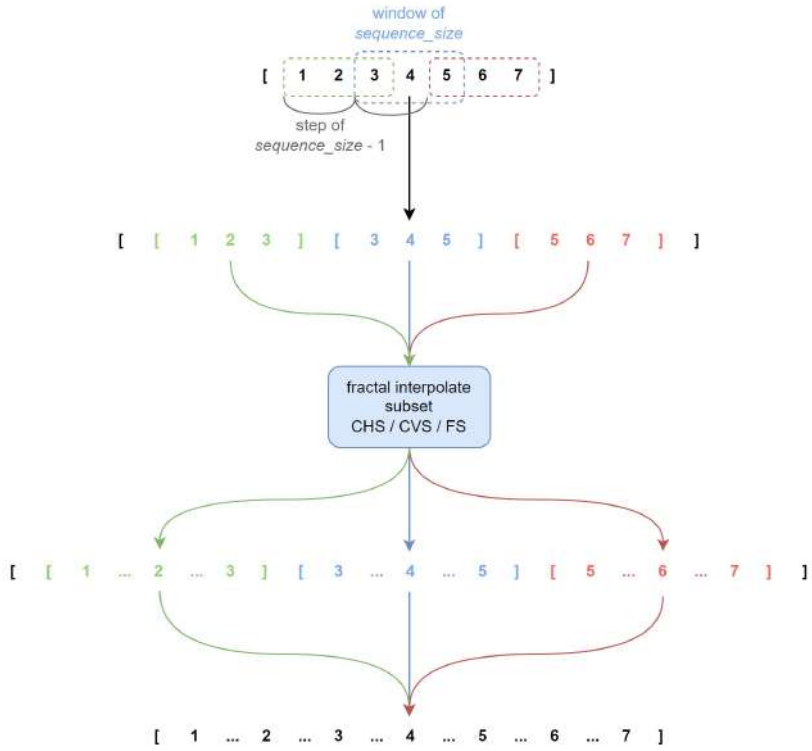

Figure 72: Interpolation flow diagram for $sequence\_size$ 3.

Algorithm 3 Fractal interpolation computation.

1: **procedure** FRACTAL_INTERPOLATION($subset$, s_n, $n\_interpolation = 17$)
2: Compute the interpolation factors a_n, c_n, d_n și e_n based on Equation (14)
3: Generate interpolation points based on Equation (13) until between every 2 points from the `subset` there are `n_interpolation` interpolation points
4: **end procedure**

The *Closest Hurst Strategy* is the first interpolation method explored, based on an approach similar to that of [142]. This strategy aims to generate interpolated data where the Hurst exponent remains close to the initial dataset's Hurst exponent, achieved through an iterative process (*Algorithm 2 - Pseudocode of Closest Hurst Strategy* from [17]). Experimental results indicated that the most appropriate vertical scaling factor, s_n, for CHS typically falls within the range of $[0, 0.2]$. Within this range, the interpolated data maintains a high degree of visual similarity to the original, with limited deviations. However,

a significant finding was that while CHS successfully preserves the Hurst exponent, it does not guarantee the preservation of other essential data properties. This limitation motivated the development of subsequent interpolation strategies designed to ensure the persistence of other critical data characteristics.

The *Closest Values Strategy* is the second interpolation method explored, and it introduces an optimized procedure for fractal interpolation, using the `Optuna` framework. For each data subset (with a `sequence_size` of 10), CVS employs `Optuna` to minimize the Root Mean Square Error (RMSE) between the fractal-interpolated subset and a linear interpolation of the original subset. This optimization process runs for 15 trials to find the optimal vertical scaling factor, s_n. The detailed pseudocode for CVS, including its objective function, is provided in Algorithm 4. It is worth noting that the `n_interpolation` parameter is assigned a default value of 17. This specific value was chosen to facilitate a conclusive comparison with the results obtained in the foundational study that provides the groundwork for this research, namely [142].

Algorithm 4 *Closest Values Strategy*.

procedure CLOSEST_VALUES_STRATEGY_OBJECTIVE($optuna\_trial$, $subset$, $n\_interpolation = 17$)
  Generate $s_n \in [-1, 1]$, a vector with the same value on all positions, representing the constant vertical scaling factor for the current `subset` in the current $optuna\_trial$ using $suggest$ API
  $interpolated\_subset \leftarrow$ FRACTAL_INTERPOLATION($subset$, s_n, $n\_interpolation$)
  $linear\_interpolated\_subset \leftarrow$ LINEAR_INTERPOLATION($subset$)
  return RMSE($interpolated\_subset$, $linear\_interpolated\_subset$)
end procedure ▷
procedure CLOSEST_VALUES_STRATEGY($subset$, $n\_interpolation = 17$)
  Create $optuna\_study$, a study with direction `minimize`, the objective function CLOSEST_VALUES_STRATEGY_OBJECTIVE() and 15 trials
  $s_n \leftarrow$ best trial parameter of $optuna\_study$
  return FRACTAL_INTERPOLATION($subset$, s_n, $n\_interpolation$)
end procedure

The results of applying CVS to all five datasets are presented in Fig. 73 - Fig. 77.

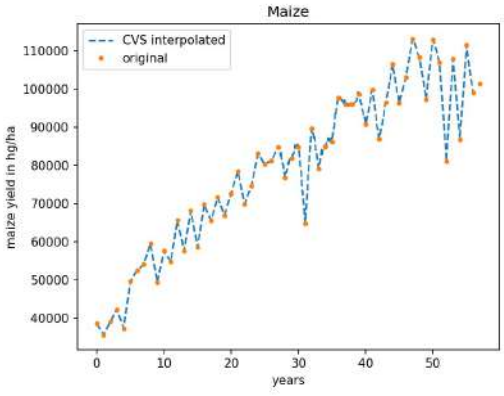

Figure 73: Maize Data set, CVS.

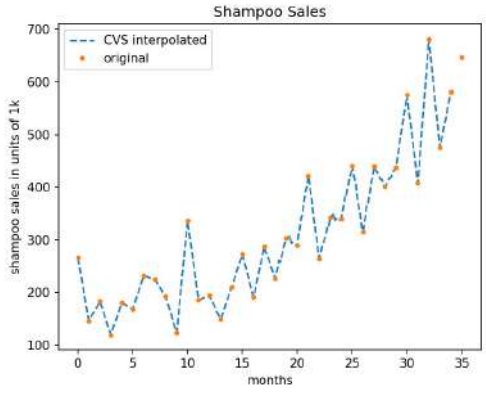

Figure 74: Shampoo Sales Data set, CVS.

A key observation from the results is that the visual appearance of the interpolated data using CVS closely resembles the results obtained with CHS when its s_n parameter was restricted to the range $[0, 0.2]$. This similarity arises because the `Optuna` optimization frequently finds the minimum RMSE for s_n values within or close to this interval. This outcome serves as a strong validation for the empirically chosen parameter range for s_n

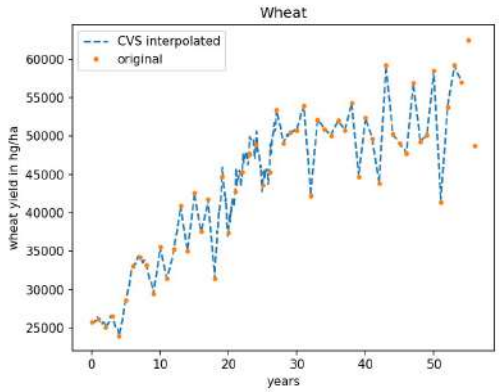

Figure 75: Wheat Data set, CVS.

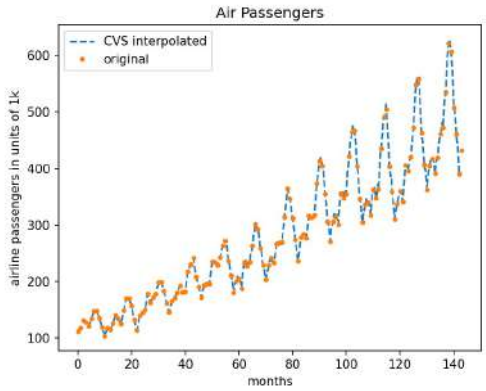

Figure 76: Air Passengers Data set, CVS.

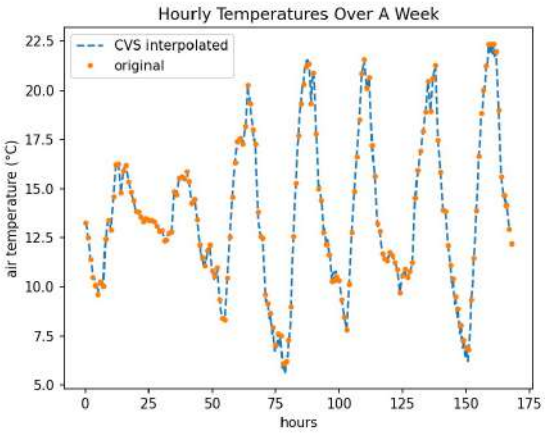

Figure 77: Weather Data set, CVS.

in the earlier CHS strategy, reinforcing its effectiveness. The evolution of the s_n parameter during the `Optuna` optimization process, highlighting its convergence, is presented in Fig. 78.

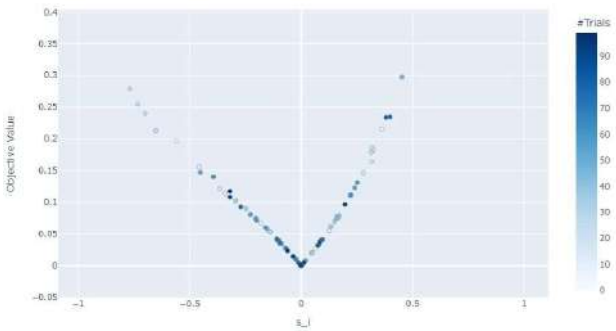

Figure 78: Evolution of parameter s_n with `Optuna`.

The third method, the *Formula Strategy* (FS), takes a distinct approach by optimizing the s_n vertical scaling factor, identifying a formula to determine a geometric method for calculating the parameter s_n. In this case, the parameter s_n is no longer constant for the entire data subset, but is obtained by applying it for each $n \in \{1, 2, \ldots, N\}$, where $N =$ `subset_size`:

$$s_n = \frac{y_n - y_{n-1}}{\sqrt{(y_N - y_0)^2 + (y_n - y_{n-1})^2}}, \tag{15}$$

However, with respect to Equation (15), a special case is observed where the denominator tends to 0. This occurs when the endpoints of the subset record similar values (i.e., the line connecting the start point to the end point is parallel to the Ox axis). This phenomenon causes the calculated value for s_n to become undesirably large in absolute value, leading to a pronounced variation in the resulting interpolation. To reduce this effect, the strategy introduces an optimization of the `sequence_size` parameter for the entire dataset, ensuring that the difference between the subset's endpoints is not negligible. This optimization is performed using the `Optuna` framework over 50 trials, searching for an optimal `sequence_size` within the range $[4, \texttt{length(dataset)} - 3]$ (Algorithm 5). This global optimization contrasts with the subset-level optimization of previous strategies.

The main advantage of FS is that, once the optimal `sequence_size` is determined for a dataset, the repetitive optimization routine required by CHS and CVS for each subset is eliminated, making it more efficient. The pseudocode for the `formula_strategy` itself is presented in Algorithm 6, noting that s_n varies for each interval within a subset, unlike the constant s_n in CHS and CVS.

Algorithm 5 Pseudocode of `sequence_size` optimization.

procedure OPTIMIZE_SUBSET_LENGTH_OBJECTIVE($optuna\_trial$, $dataset$, $n\_interpolation = 17$)
  Generate subset length $sequence\_size \in [4, \text{length}(dataset) - 3]$ in the current $optuna\_trial$ using $suggest$ API
  Split $dataset$ into $subsets$ of length $sequence\_size$
  $total\_RMSE \leftarrow 0$
  for each $subset$ **in** $subsets$ **do**
    $interpolated\_subset \leftarrow$ FORMULA_STRATEGY($subset$, $n\_interpolation$)
    $linear\_interpolated\_subset \leftarrow$ LINEAR_INTERPOLATION($subset$)
    $total\_RMSE \leftarrow total\_RMSE +$ RMSE($interpolated\_subset$, $linear\_interpolated\_subset$)
  end for
  return $total\_RMSE$
end procedure

▷

procedure OPTIMIZE_SUBSET_LENGTH($dataset$, $n\_interpolation = 17$)
  Create $optuna\_study$, a study with direction `minimize`, the objective function OPTIMIZE_SUBSET_LENGTH_OBJECTIVE() and 50 trials
  $sequence\_size \leftarrow$ best trial parameter of $optuna\_study$
  return $sequence\_size$
end procedure

Algorithm 6 *Formula Strategy*.

procedure FORMULA_STRATEGY($subset$, $n\_interpolation = 17$)
  Compute s_n based on Equation (15)
  return FRACTAL_INTERPOLATION($subset$, s_n, $n\_interpolation$)
end procedure

The optimal `sequence_size` values determined for each dataset are presented in Table 26.

Applying Equation (15) with these optimal `sequence_size` values, the interpolation results for all datasets are visually represented in Fig. 79 - Fig. 83.

With these three proposed strategies, *Closest Hurst Strategy*, *Closest Values Strategy*, and *Formula Strategy*, a synchronized comparison was performed to better under-

Table 26: Optimal `sequence_size` values for FS interpolation

Data set	*sequence_size*
Maize	29
Shampoo Sales	10
Wheat	54
Air Passengers	141
Weather	6

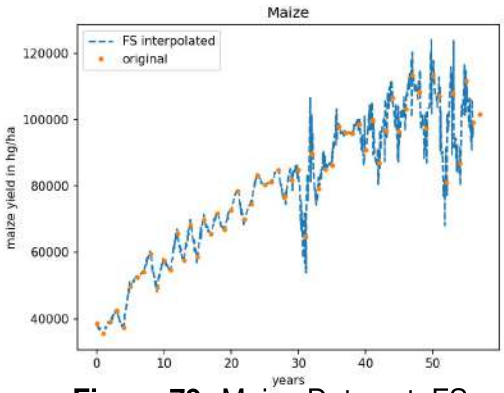

Figure 79: Maize Data set, FS.

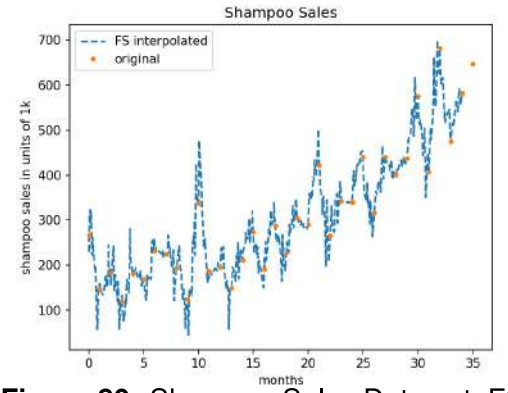

Figure 80: Shampoo Sales Data set, FS.

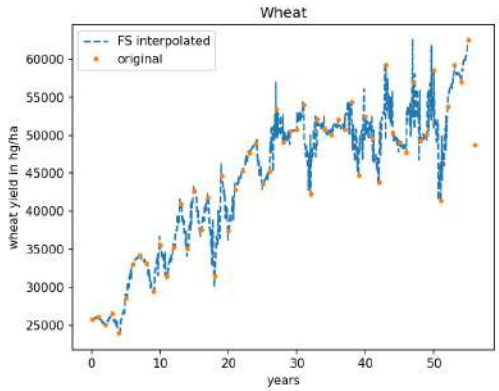

Figure 81: Wheat Data set, FS.

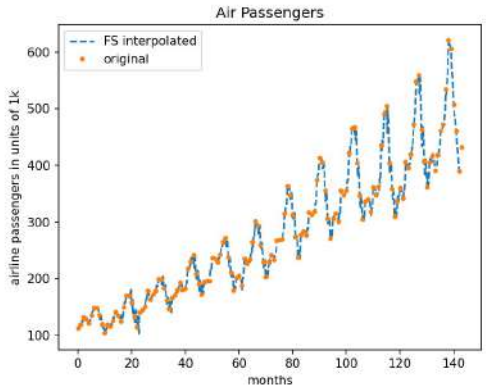

Figure 82: Air Passengers Data set, FS.

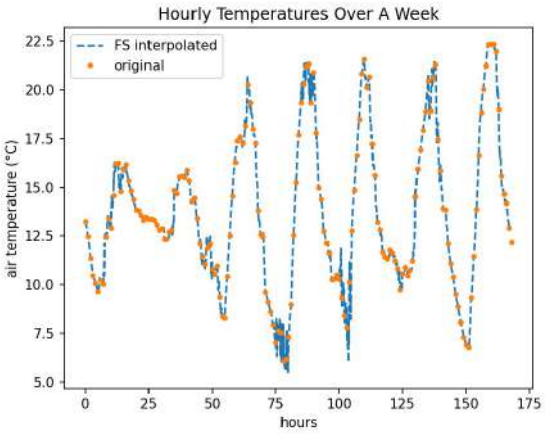

Figure 83: Weather Data set, FS.

stand their respective behaviors. In Fig. 84 and Fig. 85 are the interpolation results for Weather data set provided by the three strategies on the same graphic.

Initial visual comparisons of the interpolated data generated by the three strategies

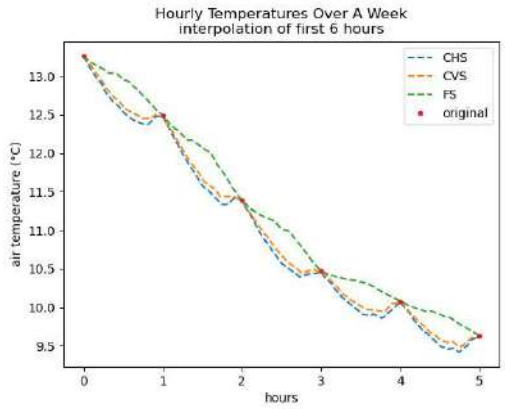

Figure 84: the first 6 points.

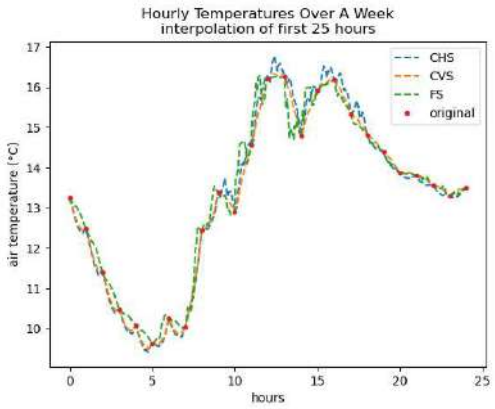

Figure 85: The first 25 points.

on the Weather dataset revealed distinct characteristics. The *Closest Hurst Strategy*, particularly when its vertical scaling factor was constrained to the $[0, 0.2]$ range, produced results visually similar to those of the *Closest Values Strategy*. In contrast, the *Formula Strategy* consistently introduced slightly higher variations in the interpolated data.

To provide a more rigorous, quantitative assessment of their performance, a specific experiment was conducted using the original Weather dataset. Hourly data points were extracted and then subjected to fractal interpolation using each of the three strategies, generating five intermediate points (`n_interpolation` = 5) to simulate 10-minute data. The Mean Absolute Error (MAE) was then computed for each strategy, providing a measure of the average difference, in degrees, between the real and the interpolated data.

For the Weather dataset, the calculated MAE values were: 0.3510 for CHS, 0.2962 for CVS, and 0.4997 for FS. These results clearly indicate that the *Closest Values Strategy* achieved the lowest MAE, thereby demonstrating its superior accuracy in this comparative test. CHS followed as the second most accurate, while FS exhibited the highest error.

Following the preprocessing stage that applies the fractal interpolation strategies described above, as outlined in the methodology in Fig. 71, the standard normalization step is performed, along with an analysis of the stationarity of the considered time series, as well as the data splitting procedure. These preprocessing steps are described in detail in the original study [17].

In the following sections, the specific structure of the employed LSTM is explored, and the final results obtained from the proposed approach are analyzed. The selected prediction model uses an LSTM layer, followed by one dense layer with dimension 1 since the datasets have only one feature.

For each data set, the number of hidden layers in LSTM was optimized using `Optuna`. For the Weather data set, the obtained values are presented in Table 27.

LSTM is configured with the default parameters. The metric used for measuring the loss in the training step is Mean Squared Error (MSE) and for the general evaluation of the model RMSE was used.

To prepare the time series data for the LSTM layer, it must be transformed into a

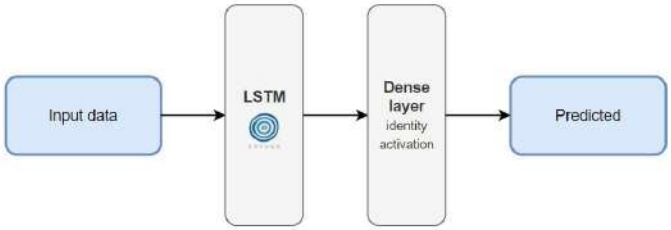

Figure 86: Neural network structure.

Table 27: Number of hidden layers

Interpolation strategy	Hidden layers
None	48
CHS	39
CVS	28
FS	10

supervised learning format. This is achieved by employing a sliding window of size `input_data_points` across the entire dataset, with a step of 1. Each resulting window serves as an input sequence for the network. For univariate time series prediction, the output corresponding to an input sequence $[data_i, \ldots, data_{\texttt{i+input\_data\_points-1}}]$ is the subsequent data point, $data_{\texttt{i+input\_data\_points}}$.

The LSTM layer in Keras (TensorFlow Keras 2.4.1) expects input data in the format of $[\texttt{batch\_size}, \texttt{input\_data\_points}, \texttt{features}]$. In this study, a `batch_size` of 1 was utilized.

An important constraint is that the `input_data_points` window size must be less than 30% of the total dataset length. Adhering to this limit is essential to ensure proper evaluation of the model's performance using a testing set, as each single-point prediction requires a full `input_data_points` sequence as input.

After completing the optimization step using the `Optuna` framework (for implementation details, refer to the original paper [17]), the configurations presented in Table 28 were obtained, along with their corresponding performance scores. The reported mean results were computed over five individual runs using the same configuration settings.

The prediction results obtained using the LSTM model on the weather dataset are illustrated, comparing the three proposed interpolation strategies with the original (non-interpolated) dataset. The visualizations highlight the improvements introduced by each strategy.

For the synthetic dataset containing daily maximum temperatures, the results are shown in Fig. 88 - Fig. 90. It can be observed that the LSTM model struggles to achieve accurate predictions on the original dataset. In contrast, the interpolated datasets yield significantly improved results, with the *Closest Hurst Strategy* delivering the best overall performance. All three strategies contribute to more accurate predictions, which can be attributed to the increased amount of training data. These findings reinforce the relevance

Table 28: Model tuning and hyperparameters optimization

Data set	Linear regression	Hidden layers	Input data Points	Epochs	Learning rate	Train RMSE	Test RMSE
Shampoo Sales	None	62	8	80	0.02	0.1717	0.4997
	CHS	48	16	15	0.04	0.0859	0.0951
	CVS	41	28	10	0.01	0.0607	**0.0749**
	FS	17	1	8	0.005	0.1218	0.1264
Air Passengers	None	56	14	65	0.007	0.0947	0.1348
	CHS	33	97	8	0.03	0.0285	0.0655
	CVS	62	51	5	0.02	0.0288	**0.0444**
	FS	29	93	7	0.03	0.0416	0.0619
Wheat	None	13	5	125	0.005	0.1480	0.2410
	CHS	8	93	12	0.03	0.0559	0.0653
	CVS	11	98	10	0.01	0.0431	**0.0552**
	FS	51	99	12	0.01	0.0656	0.0932
Maize	None	52	3	150	0.05	0.1327	0.1766
	CHS	32	82	15	0.01	0.0299	0.0304
	CVS	39	1	10	0.002	0.0187	**0.0273**
	FS	27	20	20	0.007	0.0920	0.0991
Hourly Temperatures Over A Week	None	48	8	60	0.008	0.1111	0.1383
	CHS	39	22	7	0.005	0.0325	0.0414
	CVS	28	34	15	0.001	0.3446	**0.0389**
	FS	10	17	14	0.005	0.0391	0.0462
Daily Max Temperatures	None	43	5	175	0.02	0.1132	0.4944
	CHS	47	84	12	0.03	0.0321	**0.0523**
	CVS	49	11	20	0.005	0.0290	0.0692
	FS	50	27	15	0.005	0.0526	0.1351

and effectiveness of the proposed preprocessing strategies. Additional scenarios and extended results based on the proposed methodologies are presented in [17].

]

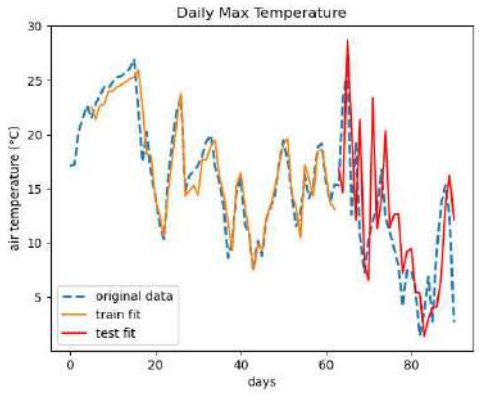

Figure 87: Prediction for daily maximum data without interpolation.

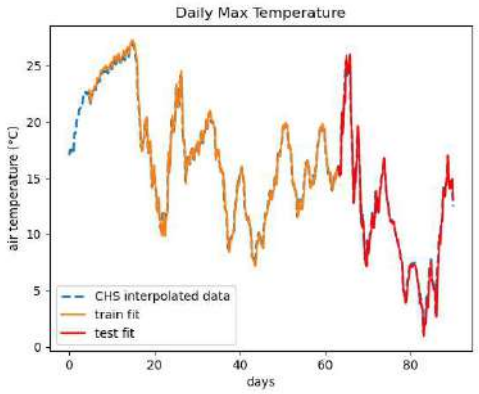

Figure 88: Prediction for daily maximum data with CHS.

In conclusion, for the majority of the datasets considered, the best results were obtained using the CVS. However, the average RMSE values recorded for CHS and CVS are similar, while the errors associated with the *Formula Strategy* are slightly higher. Nevertheless, the FS strategy remains a reliable option in the context of data generation due to the accessibility of the concepts required for its implementation. Furthermore, it is noteworthy that all results, regardless of the strategy used, are superior to the baseline case.

Following the results obtained from this study, the relevance and applications of fractal interpolation as a time series augmentation technique are confirmed and extended. The three proposed strategies thus address the optimization of the vertical scaling factor

]

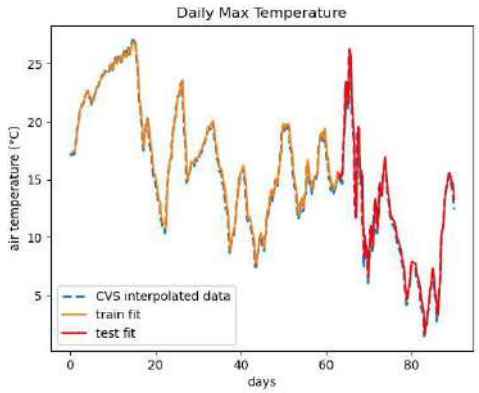

Figure 89: Prediction for daily maximum data with CVS.

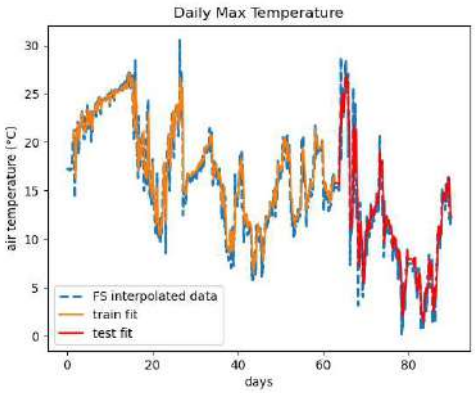

Figure 90: Prediction for daily maximum data with FS.

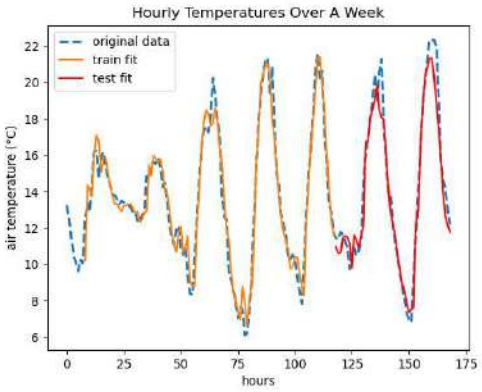

Figure 91: Prediction for hourly data without interpolation.

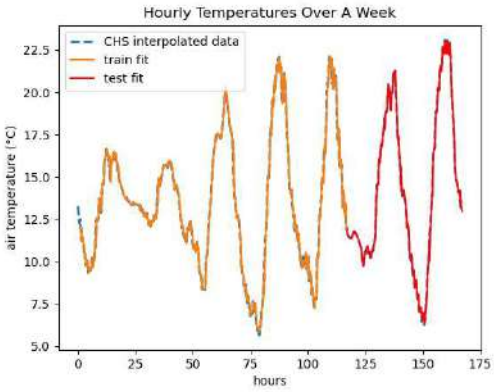

Figure 92: Prediction for hourly data without CHS.

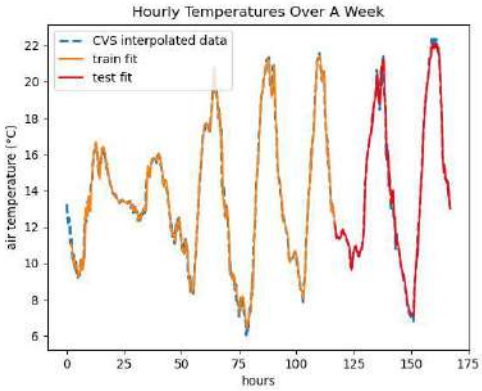

Figure 93: Prediction for hourly data with CVS.

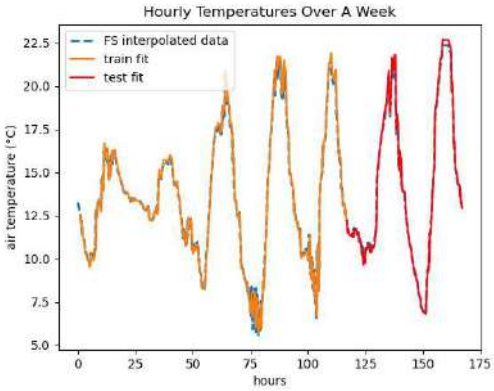

Figure 94: Prediction for hourly data with FS.

or the interpolation subset size, with the aim of generating relevant data within the problem context. Depending on the data's domain and pattern, the appropriate interpolation strategy was identified to improve predictions. Consequently, all datasets used showed an improvement ranging from 50% to 89% compared to the baseline case. All detailed results can be found in the published article coauthored by the author of this habilitation

thesis, *Fractal interpolation in the context of prediction accuracy optimization* [17], which appeared in the Q1-ranked journal *Engineering Applications of Artificial Intelligence*.

V. Digital innovation and optimization in maritime applications

The maritime sector, vital to global trade and transport, is currently facing a wave of unprecedented challenges and growing regulatory pressure. Driven by the International Maritime Organization (IMO)'s increased regulations [96], the industry faces critical demands concerning environmental impact, fuel efficiency, safety, and overall operational efficiency. Traditional maritime practices and technologies are often insufficient to meet these growing requirements for greener, safer, and more efficient operations. This critical moment offers a unique chance to adopt advanced digital technologies to modernize and optimize maritime operations.

This transformative era is characterized by a strong imperative for sustainability and technological advancement. Initiatives, such as the NEMOSHIP project, underscore the ambition to develop and deploy innovative solutions like large Battery Energy Storage Systems (BESS) and their optimal exploitation through advanced digital platforms. These platforms are redefining the maritime landscape by integrating cutting-edge digital technologies, aligning directly with the broader theme described in this chapter.

At the core of this revolution are technologies such as the Internet of Things (IoT) for real-time data collection, big data analytics for extracting actionable insights, cloud computing for scalable processing, and crucially, artificial intelligence for predictive analytics and automated decision-making. These digital solutions provide a comprehensive view of a ship's energy systems, enabling smarter decision-making and helping to optimize both performance and energy use. Examples range from Fleet Management Systems (FMS) and Voyage Optimization Platforms to sophisticated Energy Management Systems (EMS) and the emerging use of digital twin technology.

These digital innovations are having a significant impact across various aspects of maritime operations. They enhance operational efficiency by providing real-time insights and optimization recommendations, leading to reduced fuel consumption, lower operating costs, and improved vessel performance. Furthermore, they significantly contribute to sustainability by lowering greenhouse gas emissions and assisting the industry in meeting environmental regulations.

To concretely illustrate the power and applicability of the digital innovation and optimization strategies discussed, and to showcase their practical implementation within the maritime domain, this section will explore two distinct case studies. These studies, [21] and [18], developed through collaborative research efforts under the NEMOSHIP project, directly address the aforementioned challenges of enhancing efficiency, sustainability, and safety in complex maritime systems through advanced monitoring, diagnosis, and optimization. Central to both of these investigations is the *NEMOSHIP Digital Cloud Platform*. This proposed platform is built on principles of distributed computing and a mi-

croservices architecture, ensuring a modular design, Fig. 95. This modularity enhances flexibility, scalability, and simplifies maintenance by allowing independent modification or replacement of individual components. Given its cloud-deployed nature, effective inter-module and external communication relies on robust network-based protocols, critical for efficiency, security, and reliability.

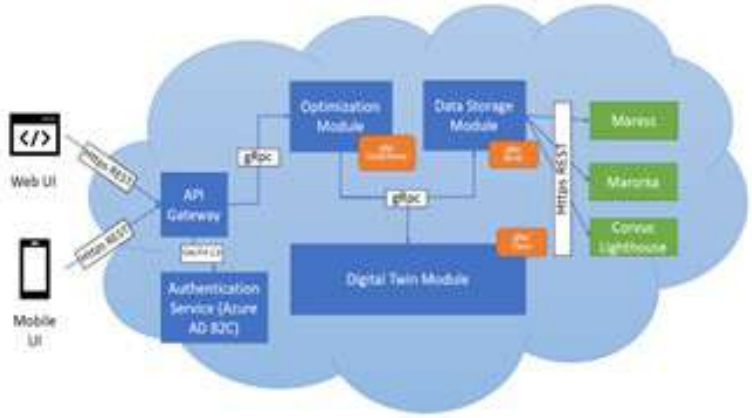

Figure 95: Module-to-Module communication.

The platform integrates several key modules:

A. Digital Twin module - This module is responsible for highly accurate simulations of real-world ships, using state-of-the-art simulation software such as Simcenter Amesim (Fig. 96).

B. Optimization module - Composed of two main services, the Optimization Management Service and the Optimization Engine Service, this module processes user input to find optimized values for variable parameters, providing crew guidance data (Fig. 97). The Management Service handles user inputs via an API, while the Engine Service focuses on calculating optimized values.

C. Data Storage module - Central to data handling, the Data Storage module's core component is the Data Adapter Service. This service exposes a gRPC interface (Data Adapter Service API) for inter-module communication. It includes subcomponent adapters for various data providers, managing data retrieval, formatting, normalization, and correlation before passing it back to the client.

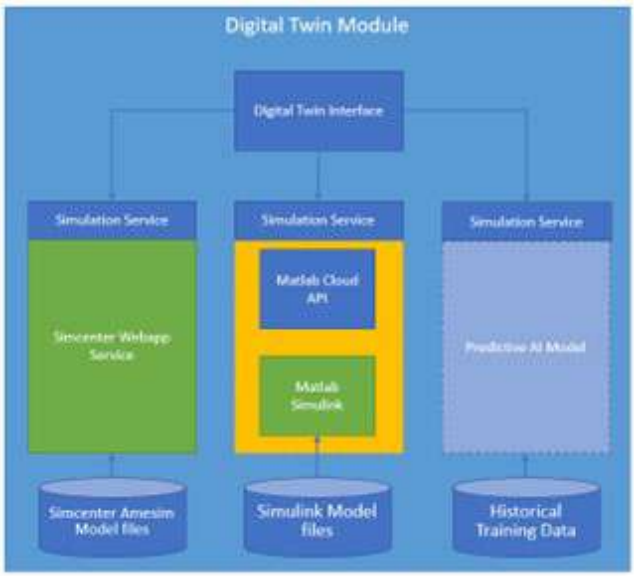

Figure 96: Digital twin module.

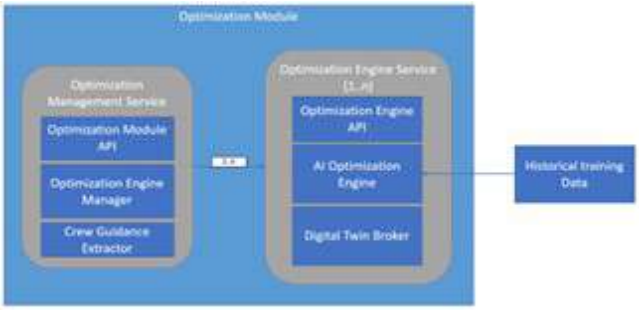

Figure 97: Optimization module.

Corelated with [C.] key module, the comprehensive data managed by this module is instrumental in enabling advanced historical analysis and "what-if" scenario exploration. This includes enabling the comparison of two simulation results for the same historical voyage: one using default parameters mirroring actual operational conditions, and a second allowing for user-defined adjustments to operational parameters (such as fuel type, genset order and load-dependent start/stop, or Battery Management System settings like peak shaving, zero emissions, spinning reserve, min/max state of charge, and charging limits) and constructive parameters (like battery capacity). These simulations are then validated against real-world data (e.g., from external APIs like Marorka). For instance, for fuel consumption analysis, the platform enables a direct comparison of the results from both simulations, validated against Marorka data, with the outcomes depicted in Fig. 98. These analyses, which uses Simcenter Amesim digital models for robust and realistic system behavior modeling, allow for performance and efficiency evaluation across custom time intervals or defined port-to-port trips. The module also supports archive functionality, allowing users to revisit past results without re-executing simulations.

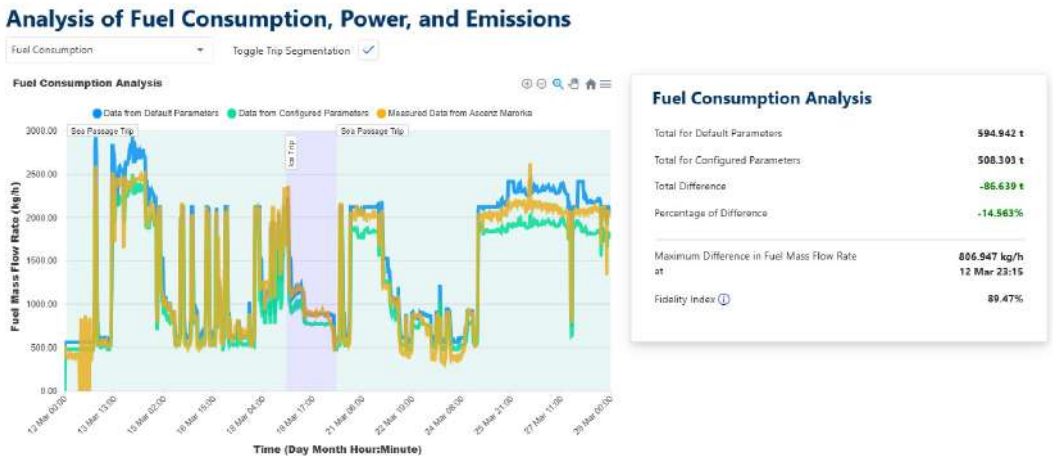

Figure 98: Historical trip analysis – Results step.

The *NEMOSHIP digital platform* integrates several key external components, which are existing commercial products or partner technologies widely used in the maritime industry. This section details a few representative examples:

- Ascenz Marorka Web API - Provides a standardized interface for secure access to Ascenz Marorka's extensive maritime data center, enabling seamless data retrieval and manipulation for external systems.
- Corvus Lighthouse - A data logger that collects, aggregates, and efficiently transfers data from Corvus Energy BESS (Battery Energy Storage System) inter-pack communications networks to cloud storage for processing.
- Simcenter Webapp Cloud Platform (powered by Amesim) - A flexible tool that makes system simulation models accessible to everyone through a web-based platform. It enhances accessibility, customization, and security for leveraging simulations across various departments and user types within an organization.

Three distinct scenarios are planned to be supported by the proposed architecture, see Fig. 99. One scenario is the *battery usage optimization during trip planning*. This scenario is specifically designed to support the crew in optimizing the Battery Energy Storage System (BESS) for upcoming voyages, with the primary objective of improving energy efficiency and reducing operational costs. Users can input or define trip parameters within the *NEMOSHIP Digital Platform* to receive optimized guidance values for vessel operation, through the integrated digital twin. Optimization is performed on a predefined route, which can be imported into the Digital Platform (e.g., as an .xlsx file). The platform further supports the definition of state changes via waypoints and optional time differentials. The underlying optimization algorithm uses as input the power profile generated by the digital model, which accounts for the specified route, weather conditions, and ice conditions. Simulation results are stored in the database, enabling download in CSV format, and the system also facilitates the comparison of two or more optimization results simultaneously.

The *NEMOSHIP Digital Platform* acts as a complementary tool to existing maritime trip planning software, not a replacement, facilitating the import of waypoints and other relevant data from established maritime systems to enhance battery usage optimization. The second scenario is *at sea battery usage optimization*. This scenario demonstrates the system's dynamic adaptability from trip planning to real-time at-sea operations. While initial trip planning involves defining waypoints and state changes, once the vessel is at sea, the platform's interface dynamically updates to reflect current conditions. It continuously tracks vessel position and intelligently filters and displays only the waypoints and state changes relevant to the current leg of the journey. This ensures the crew receives precise, real-time information, preventing overload and enabling focused decision-making for effective battery efficiency management. The last use case, *at sea scenario prediction* employs the vessel's digital twin to enable advanced "what-if" scenario simulations during at-sea operations. By manipulating various input parameters, environmental conditions, and operational modes within the digital twin, the platform can forecast vessel performance, including fuel efficiency, speed, and maneuverability. This capability allows for comparative analysis of parallel simulations, providing critical insights into how different operational strategies impact key performance indicators, thereby supporting strategic decision-making and optimizing for efficiency and sustainability.

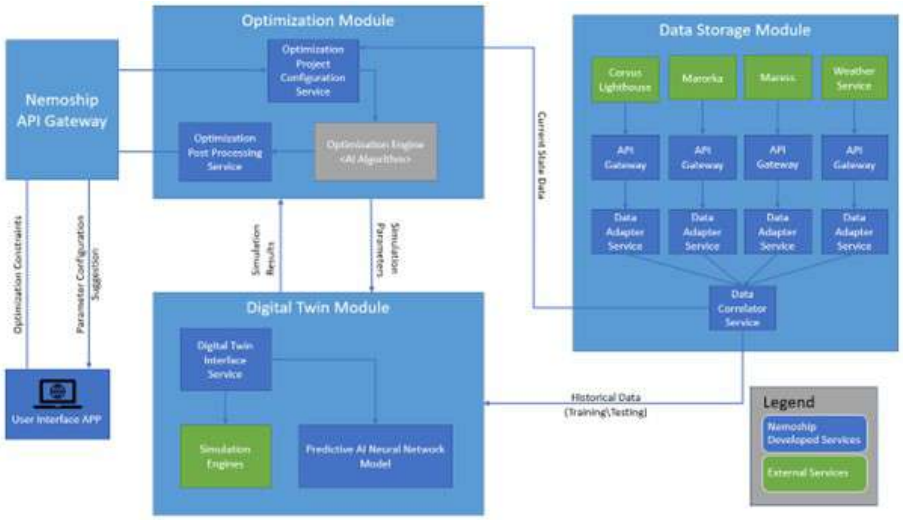

Figure 99: An overview of the NEMOSHIP proposed architecture.

In conclusion, the convergence of automotive and maritime technologies is fundamentally transforming the maritime sector, opening up unprecedented opportunities for enhancing efficiency, sustainability, and safety. By leveraging advancements in areas such as electric propulsion, advanced sensor networks, and smart connectivity solutions, the maritime industry can address long-standing challenges and drive transformative change. This cross-industry innovation promises a future where operational efficiency is maximized, environmental impact is minimized, and safety standards are elevated.

Building on the information presented in this section, the included case studies serve as concrete examples of this cross-industry innovation, demonstrating practical applica-

tions of digital innovation and optimization strategies within the maritime domain. The papers [21] and [18] represent an initial, yet significant, step in harnessing this synergy to advance monitoring, diagnosis, and optimization in complex maritime systems. These contributions have been recognized through their peer-reviewed publication and presentation at Scopus-indexed conferences: the *2024 IEEE International Conference on Electrical Systems for Aircraft, Railway, Ship Propulsion and Road Vehicles & International Transportation Electrification Conference* (ESARS-ITEC) and the *CONAT 2024 International Congress of Automotive and Transport Engineering*. This marks a promising beginning in charting a course towards a greener, smarter, and more sustainable maritime future. These works have received funding from the European Union's Horizon Europe Research and Innovation programme under grant agreement No 101096324 (NEMOSHIP project).

Chapter 3. Computational models and intelligent behaviors

Problem statement and research context

This chapter explores how neural cellular automata (NCA) and cellular automata (CA) serve as advanced computational models, capable of exhibiting intelligent behaviors through pattern generation, texture synthesis, and predictive analytics. My interest in this field began in 2008, with the introduction of formal languages and compilers in my teaching, a field where finite automata (FA) are a foundational concept. Approximately four years ago, the curriculum of this course has been expanded to include more advanced projects aimed at connecting theoretical computer science with recent developments in AI. These efforts explored the relationships between neural networks and key elements of formal language theory, such as automata, grammars, and parsing mechanisms. This exploration eventually introduced me to a relatively recent yet remarkably powerful concept: the cellular automaton, and in particular, the neural cellular automaton.

Despite their differences, FA and NCA share foundational similarities as discrete computational models. Both operate with finite states, evolve in discrete time steps, and are governed by well-defined local rules that determine state transitions based on local information (e.g., input for FA, neighbors for NCA). They are rule-based systems capable of generating complex emergent behaviors from simple components. While both FA and NCA rely on local rules for state evolution, their architectures and capabilities diverge significantly. Finite automata are abstract computational models characterized by a finite set of states and transitions between them, designed to process strings over a finite alphabet and to recognize regular languages. In contrast, NCAs are grid-based computational systems in which each cell updates its state using a shared, trainable neural network. These systems learn complex behaviors through gradient-based optimization on data-driven tasks. This enables NCAs to excel in tasks such as image regeneration, pattern formation, and morphological computation, effectively bridging traditional cellular

automata with modern ML paradigms.

The sustained focus on NCAs, driven by the aim to understand their unique capabilities and apply them to specific computational problems, naturally led to a series of dedicated research efforts. Below, I present the principal studies whose findings support this research direction.

Highlighted studies and key findings:

- **A. Băicoianu**, R. Ivan, I. Popa, 2025 - *Exploiting cellular automata and linear regression to predict disease spread* [11]

  Main findings:

  - A model was developed and applied, combining non-uniform cellular automata with linear regression and the SIRD epidemiological model for predicting Covid-19 spread.
  - The model was successfully adapted to assess pandemic progression in Germany and Italy, incorporating the impact of neighboring countries.
  - A new strategy, to expand prediction to countries lacking infection data (e.g., Switzerland) by using estimated parameters from neighboring regions and randomly initialized parameters, was introduced.
  - The model demonstrated high precision in tracking infection rates over time and was validated by reliable public data sources.
  - The work provided valuable insights into the regional distribution of the epidemic's impact and significantly contributed by extending the model's application beyond single-country borders to extrapolate infection patterns across national boundaries.

- S. Catrina, M. Catrina, **A. Băicoianu**, I. C. Plajer, 2024 - *Learning About Growing Neural Cellular Automata* [33]

  Main findings:

  - The paper explored how different image characteristics influence the development and stabilisation processes of a 2D NCA and proves how certain features can influence NCA learning and evolution over multiple time steps.
  - One objective of this paper was to understand how the NCA responds to inherent characteristics of the training image, such as symmetry, color and complexity. Taking into consideration how the NCA learns to evolve, multiple hypothesis have been raised in order to explain why certain behaviours are observed. The main experiment in this regard does not use a stabilizing strategy, and highlights how non-symmetric images are harder to reproduce, how border affects the way the NCA evolves, how diagonal lines seem to appear faster and more wide-spread.

  - Another experiment explored how a pooling-based training strategy affects the evolution of the NCA and its stabilization over time. Here the finding that such pooling strategy stabilizes the evolution of the model was reproduced, as pointed out by previous studies.
  - A secondary objective of the study was to explore the effects of different filters on the training strategy and consequently on the evolution of the automaton. Therefore, apart from the Sobel filters previously used, experiments were conducted using their rotated variants, using only the Laplace filter and combining the Laplace with the Sobel filters. The conclusion of these experiments was that it may be possible that by integrating multiple filters, such as the Laplace filter, the NCA may be able to learn a more reliable way of evolution towards the target image.

- M. Catrina, I. C. Plajer, **A. Băicoianu**, 2024 - *Multi-Texture Synthesis through Signal Responsive Neural Cellular Automata* [31]

  Main findings:

  - The study proposed a single NCA architecture capable of generating multiple distinct textures from a unified model through genomic coding, effectively addressing the scalability limitations of previous single-texture NCAs. The proposed architecture, with less than 10k parameters, was trained to synthesize diverse batches of textures, including similar, dissimilar, structured, and non-structured examples, and consistently achieved high-quality results across all scenarios. The solution achieves a Learned Perceptual Image Patch Similarity (LPIPS) score below 0.45, a competitive result considering the model's lightweight architecture, simple perceptual design, and the inherent visual complexity of irregular textures.
  - The work demonstrated that the trained NCA preserves its regenerative abilities while enabling smooth interpolation between texture encodings, leading to the synthesis of hybrid textures. Furthermore, it supports editing operations like grafting, allowing multiple distinct textures to coexist within the same grid, underscoring the approach's spatial composability and flexibility.
  - The study highlighted the critical influence of genome encoding and loss function design on the training dynamics and convergence behavior of the model, directly affecting texture fidelity and generalization across a range of texture types, including regular, irregular, and microtextures. It is shown that, within the NCA architecture, Sliced-Wasserstein Loss outperforms in handling irregular patterns, while OTT Loss yields better results for regular textures. Both approaches achieve a Gram Matrix Distance of approximately 0.01 from the original textures, indicating high perceptual fidelity.

All these research contributions have advanced the field of computational models and intelligent behaviors, with each study being disseminated through recognized academic publications, specifically published in C-ranking [11], B-ranking [33], and A-ranking [31] journals, respectively.

In the following sections, in-depth analyses of two key studies are presented, highlighting the main research trends in applying NCA for generative design and complex pattern formation. First, the research exploring how NCAs acquire the ability to grow and regenerate specific shapes is examined, with an analysis of the factors influencing their training and stability [33]. Second, the work demonstrating how a single NCA can learn to synthesize multiple textures and exhibit advanced features such as interpolation and regenerative capabilities is discussed [31]. These studies serve as excellent examples, chosen for their significant contributions to understanding the fundamental mechanisms of NCA learning and their innovative applications in self-organization.

I. Understanding learning and growth dynamics in NCAs

Problem statement

Cellular automata are discrete computational models developed to simulate cell behavior and emulate properties of real-world organisms like local interaction, parallelism, and self-replication [85]. They are powerful for modeling complex systems through simple local update rules, especially when combined with learning automata [88]. A CA consists of a grid of cells, each with a state that updates at discrete time steps based on a predefined rule, which considers its current state and the states of its neighbors. CAs can be broadly categorized into regular cellular automata, with manually defined rules, and NCA, which utilize neural network-based update rules [88], [73].

The earliest forms were elementary cellular automata, one-dimensional systems with two-state cells [182]. Their evolution is determined by a rule based on a cell's state and its immediate neighbors, leading to 2^8 possible rules. Despite their apparent simplicity, elementary cellular automata have been extensively studied for their diverse and complex emergent behaviors [185], [26]. As the field advanced, interest shifted to bidimensional CAs, where cell updates depend on neighborhood types like Moore (8 neighbors) or Von Neumann (4 neighbors) [85], among others [193], [28]. The pivotal introduction of neural networks into cellular automata led to NCAs, enabling the learning of complex update rules. This innovation has opened new directions for applications such as texture generation and regeneration [116], [31], and the growth and regeneration of 2D or 3D models [76], [118] tackling problems previously impossible with fixed rules.

This particular section explores the behavior of an NCA specifically designed for growing and maintaining a 2D model. Building upon existing NCA architectures [118], this research analyzes the evolution and learning process of such an NCA when trained on various images. This approach investigates how different image characteristics influence its development and compare the resulting behaviors. Furthermore, a key aspect of this

work involves analyzing the performance of a stabilized model, which successfully maintains its shape after growth, addressing a limitation of earlier models. The experiments were motivated by the desire to understand precisely what and how such an automaton learns during training, and to identify the factors that contribute to an image's susceptibility to self-destruction or shape its evolution during NCA expansion.

Methodology and experimental insights for NCA

Motivated by the methodology from [118], an NCA designed for image generation was tested. The update rule, which determines how each cell changes depending on its neighbors, will now be learned by a neural network. Prior to examining the neural network, the steps involved in NCA training will be reviewed in detail, with a focus on cellular automata.

In this NCA, a cell's state, traditionally binary (dead/alive), is represented by a continuous vector of 16 values, enabling better compatibility with neural networks. The first four values encode pixel information: positions [1-3] hold normalized RGB values (0-1), and position 4 contains the alpha channel (α). The alpha channel is important for determining a cell's status: $\alpha = 0$ signifies a dead cell, $\alpha = 1$ a fully developed cell, and values $\alpha \geq 0.1$ (threshold l_s) indicate a growing state. The remaining 12 values (positions [5-16]) are additional channels used by the neural network to develop complex local rules.

Initially, all pixels except the seed pixel have all 16 feature values set to 0. The seed pixel is initialized with its target RGB values (normalized), an alpha channel of 1, and the remaining channels set to 0. Fig. 100 illustrates how these feature values evolve within a cell over iterations.

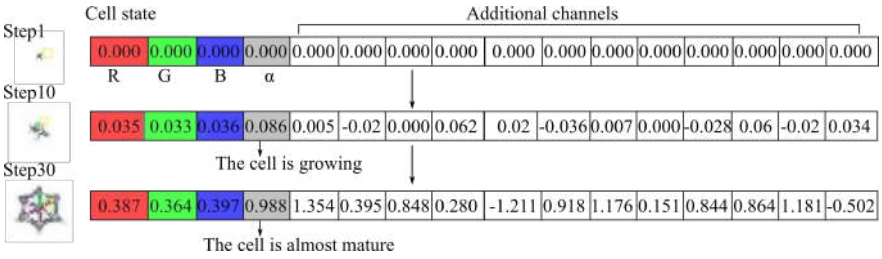

Figure 100: The encoding strategy.

The update rule is one of the main mechanisms governing the dynamic behavior of any cellular automaton. This rule takes as input a cell's current state and its neighbors' states, outputting an "update value" that dictates how much the cell's features and state modify. Significantly, these updates are applied simultaneously across the entire grid, as the image is processed through the neural network at once, enabling the image to evolve by iterating this process.

The first stage of the neural network involves convolving classical 3×3 filter masks with each channel of the input image, operating on the Moore Neighborhood. These filters are essential for extracting local spatial information, which serves as the primary input for the subsequent neural network layers. The filter masks considered are:

- $Sobel_x$: A high-pass filter that emphasizes vertical edges by considering 6 neighboring cells.
- $Sobel_y$: A high-pass filter that emphasizes horizontal edges by considering 6 top-bottom neighbors.
- $Identity$: Ensures that the current cell's value is directly incorporated and influences the update rule.

The choice of Sobel filters is motivated by their effectiveness in identifying relevant image features, particularly edges and their orientation, which are essential for the neural network to learn complex patterns, as demonstrated in [118]. In this approach, additional filters with other orientations, such as diagonal Sobel and the isotropic Laplace filter, with their results detailed in the experimental section. Each filter is applied to each channel of the input, generating a $72 \times 72 \times 16$ array. These three arrays are then concatenated to form the network's input feature array of size $72 \times 72 \times 48$. This design strictly follows the fundamental cellular automaton principle that a cell's update is based solely on local information from its neighborhood.

The core of an NCA lies in its neural network, which is responsible for learning and applying the complex update rule for each cell. The neural network's architecture, detailed in Fig. 101 and Fig. 102 consists of a perception layer followed by two convolutional layers.

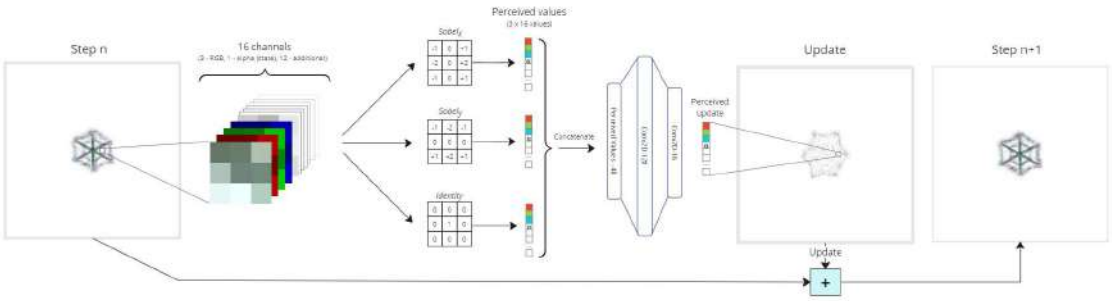

Figure 101: Overall architecture of the NCA model.

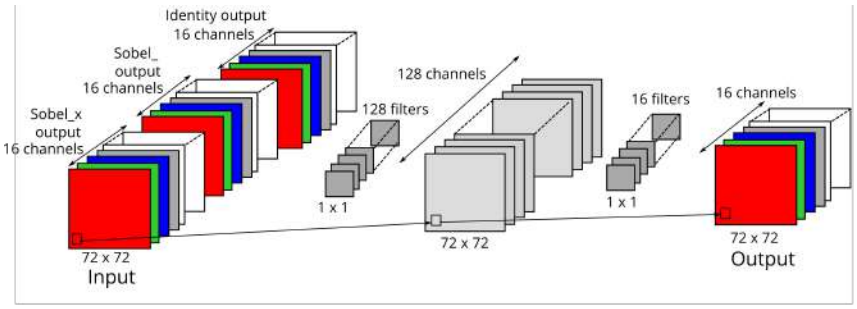

Figure 102: Detailed architecture of the neural network.

The network processes the $72 \times 72 \times 16$ input image (representing the current cell states). The initial perception layer applies 3×3 filters ($Identity$, $Sobel_x$, $Sobel_y$ to each

channel, effectively enabling each cell to perceive its 3×3 neighborhood. This results in a $72 \times 72 \times 48$ tensor. This tensor then passes through two subsequent 1×1 convolutional layers: the first with 128 kernels and ReLU activation, and the second with 16 kernels. The final output is a $72 \times 72 \times 16$-dimensional array, representing the update values that are added to the original cell states.

The application of this update rule for a single time step is synthesized in Algorithm 7. It involves selecting the living cells and their neighbors, applying the perception filters, and then passing the resulting features through the convolutional layers to obtain the update values. These values are then added to the corresponding cell states. The complexity of this single update step is proportional to $p \cdot n$, where n is the number of cells and p is the number of learnable weights (approximately 8000).

Algorithm 7 The application of the update rule.

Input: Current state of the automaton (the image)
Output: The updated state of the automaton after one step

function APPLY_UPDATE_RULE(state_nca)
    $cells\_to\_update \leftarrow get\_living\_mask(state\_nca)$
    $features\_vector \leftarrow apply\_perception\_filters(state\_nca)$
    $cell\_update\_values \leftarrow forward(features\_vector)$
    for $cell \in cells\_to\_update$ **do**
        $state\_nca[cell] \leftarrow state\_nca[cell] + cell\_update\_values[cell]$
    end for
    return $state\_nca$
end function

During training, the NCA evolves over multiple steps, typically between 64 and 96 iterations, as recommended in [118] (Algorithm 8). The neural network's weights and biases are adjusted via backpropagation after such an evolution cycle is completed. The optimization aims to minimize the L_2 loss between the RGB channels of the evolved NCA's output and the target image. This multi-step evolution provides a sufficient time interval for the neural network to learn effective update rules. The overall complexity of the evolution process is proportional to $num\_steps \cdot A1$, where $A1$ is the complexity of Algorithm 7.

Algorithm 8 Evolution of the NCA.

Input: The automaton (nca), number of steps to evolve (num_steps)
Output The evolved nca

function EVOLUTION_OF_THE_NCA(nca)
    for $steps \leftarrow 1; steps \leq num\_steps$ **do**
        $nca \leftarrow apply\_update\_rule(nca)$
    end for
    return nca
end function

After the NCA is trained, its lifetime simulation begins from a single living cell. The trained neural network then applies the update rule iteratively to all living cells and their neighbors. This process can run indefinitely until the desired output is achieved, though in the experiments it was stopped at 2000 iterations. Observations showed that images tend to distort and self-destruct after approximately 300 iterations, indicating a stability challenge addressed in later sections.

The training process for the NCA model outlined in Algorithm 9, spans 8000 steps. Each training step involves three key phases:

1. *Seed image initialization:* The process starts with a "seed" image, where all pixels are background (dead, $\alpha = 0$) except for a single central pixel which is alive ($\alpha = 1$). This seed serves as the starting point for the growth of the NCA.

2. *NCA evolution:* The NCA is allowed to "live" for a randomly selected number of iterations, between 64 and 96 steps, by repeatedly applying the update rule (as described in Algorithm 8). This allows the automaton to evolve, in order to grow the target image.

3. *Loss calculation and weight adjustment:* After the evolution, the L_2 loss is calculated by comparing the RGB channels of the evolved NCA's output with those of the target image. This loss is then used to adjust the neural network's weights and biases via backpropagation.

Algorithm 9 Training of the NCA.

Input the nca, representing the seed image, and the target image
Output the trained nca

function TRAIN_NCA(nca, target)
  for $iteration \in range(0, number\_of\_epochs)$ **do**
    $num\_steps \leftarrow random\ select\ a\ number$
    $of\ steps\ in\ range(64, 96)$
    $nca \leftarrow evolution\_of\_the\_nca(nca, num\_steps)$
    $loss \leftarrow L_2(nca, target)$
    $update\ weights\ by\ backpropagation$
  end for
  return $trained\_nca$
end function

Each NCA is trained to generate a specific image; for instance, a spider web image would serve as the target data for its dedicated NCA. The training strategy yielded compelling results, with typical loss histories resembling those shown in Fig. 103.

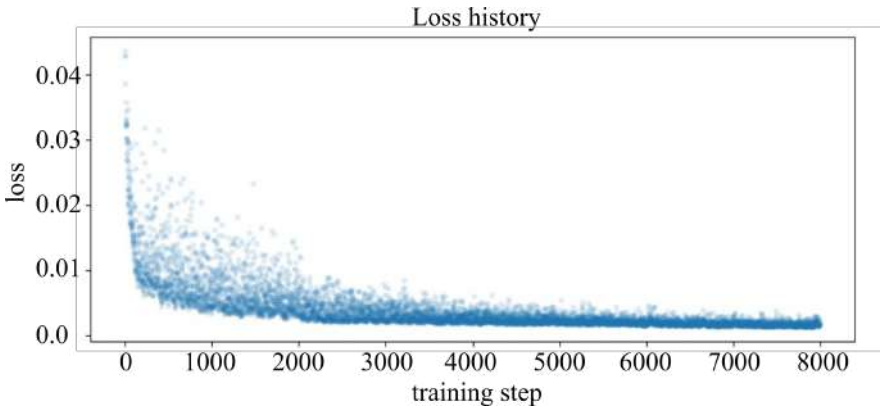

Figure 103: Loss history for spiderweb images during training.

The initial training approach, while yielding interesting results, proved unsatisfactory regarding the stability of the generated images. The automaton often continued to expand or distort even after reaching the target state, exhibiting unpredictable behaviors such as continuous growth, disappearance, or minor alterations. Achieving stable image generation was therefore a critical challenge.

To address this, a pooling strategy was employed, as presented in the original paper [118], rather than simply extending the training duration of a single iteration. This strategy, conceptually illustrated in Fig. 104, is designed to expose the neural network to a diverse range of image states during training, thereby fostering robustness and stability.

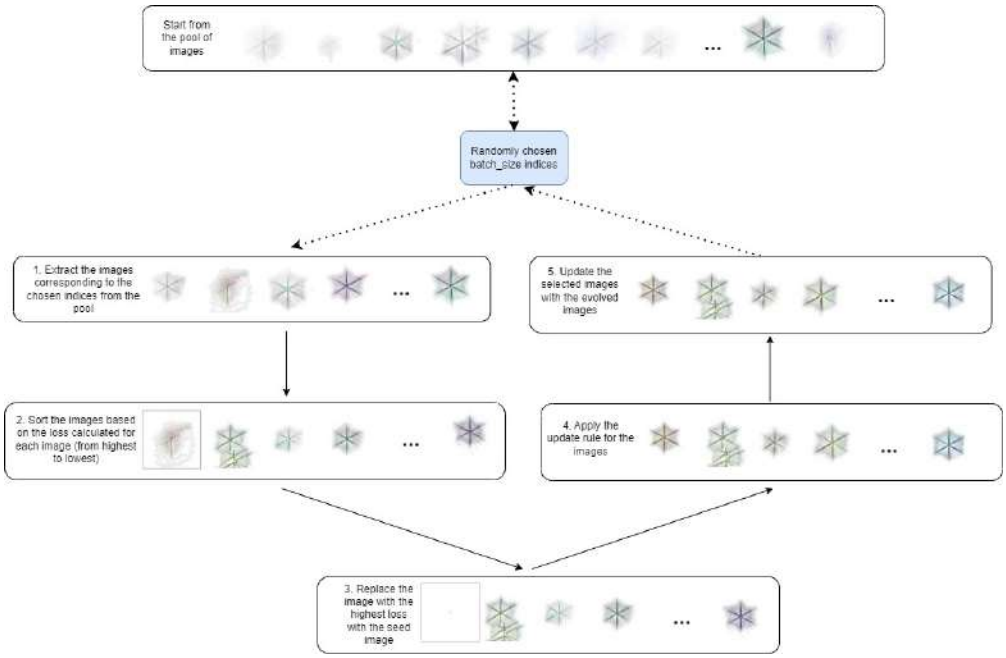

Figure 104: Pooling strategy depicted for one epoch.

The pooling strategy operates as follows [33]: A large pool of images (e.g., 1024 in these experiments) is maintained. Initially, this pool consists entirely of seed images, each of them containing a single alive cell. In each training step:

1. A batch of individuals (e.g., 8 images) is randomly selected from the pool.
2. The individual within this batch that exhibits the highest loss (i.e., the worst performer) is identified and replaced with a fresh seed image. This crucial step ensures that the neural network continuously learns to grow from a single cell and is not solely trained on already evolved or distorted images.
3. All images in the selected batch are then evolved by applying the NCA's update rule for a random number of steps (between 60 and 90 iterations).
4. The loss is calculated for the evolved batch, and the neural network's weights are updated via backpropagation.
5. Finally, the evolved images from the batch replace their original counterparts back into the main pool.

This dynamic process allows the neural network to train on images at various stages of evolution, from initial seeds to partially grown or even slightly corrupted forms, as visually demonstrated by samples from the pool in Fig. 105. This approach helps the neural

network learn how to initiate growth, maintain structure, and potentially recover from minor deviations.

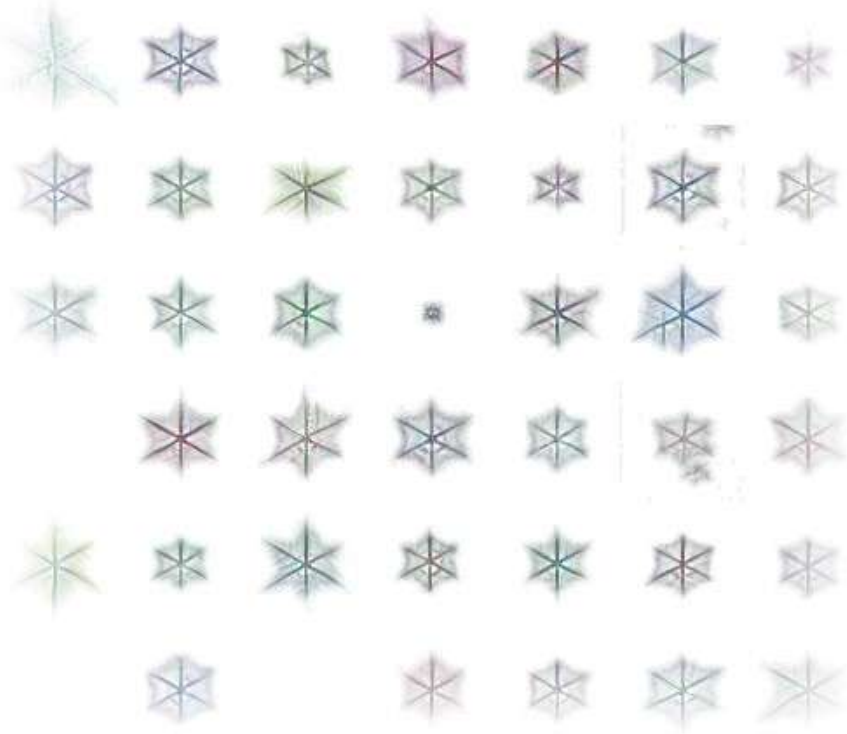

Figure 105: Pool of images observed to be in different stages of training.

This training strategy successfully led to more stable loss plots, exemplified by Fig. 106. After training, the term *evolve* refers to applying the now-fixed update rules learned by the neural network to a new single living cell, observing its behavior as it attempts to grow and maintain the target image.

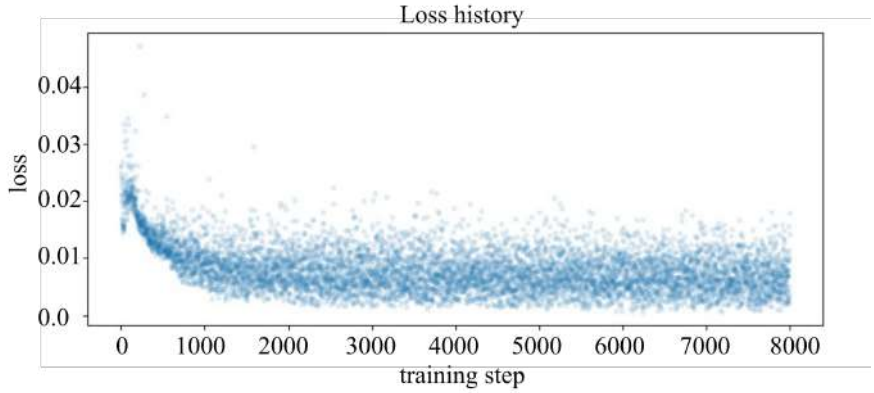

Figure 106: Loss history of spiderweb image with pooling strategy.

The research discussed in this section aimed at achieving two primary objectives: first, to understand the dynamics of NCA expansion and how inherent characteristics of target images (e.g., symmetry, orientation, contour) influence both growth and stabilization; and second, to investigate the impact of different filter orientations on the NCA's

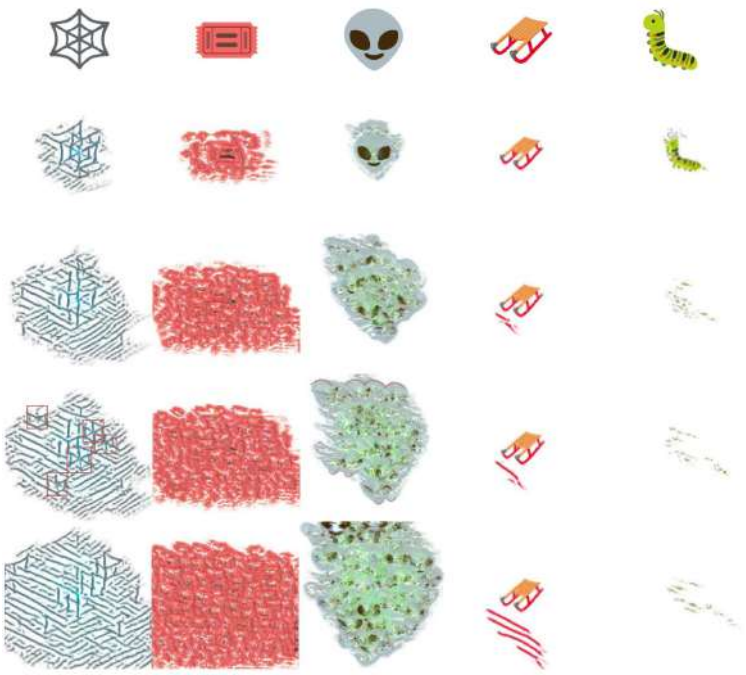

Figure 107: Expansion dynamics for figures with varying symmetry, illustrated across 300, 500, 600, and 800 iterations.

learning and growth processes. To address these objectives, three main experiments were designed:

Experiment 1: Unstabilized NCA growth. This experiment utilized the training approach described in Algorithm 9 (starting always from a seed image) without any explicit stabilization strategy. It aimed to analyze the unpredictable expansion behavior of the NCA on various target images, as exemplified in Fig. 107. The focus was on how image features like symmetry, orientation, and object contour affect the automaton's uncontrolled growth and distortion after reaching the target shape.

Experiment 2: Stabilized NCA with pooling strategy. To address the stability issues observed in Experiment 1, this experiment trained the NCA using the pooling strategy detailed in *Algorithm 4 - Pooling Strategy*, from [33] which is summarized in the steps described earlier. The goal was to enable the cellular automaton to not only reach its target image but also to cease expanding and maintain its shape thereafter.

Experiment 3: Influence of filter orientations. This experiment explored the impact of replacing or combining the initial $Sobel_x$ and $Sobel_y$ filters (used for feature selection) with other high-pass filters. The findings indicated that directional filters are crucial for consistent figure growth, while isotropic filters like Laplace alone produce unusable results. However, combining directional Sobel filters with Laplace was shown to enhance performance.

All simulations involved feeding a single image with a central alive cell to the trained NCA, allowing it to evolve over a specified number of steps, with each step representing one pass through the network.

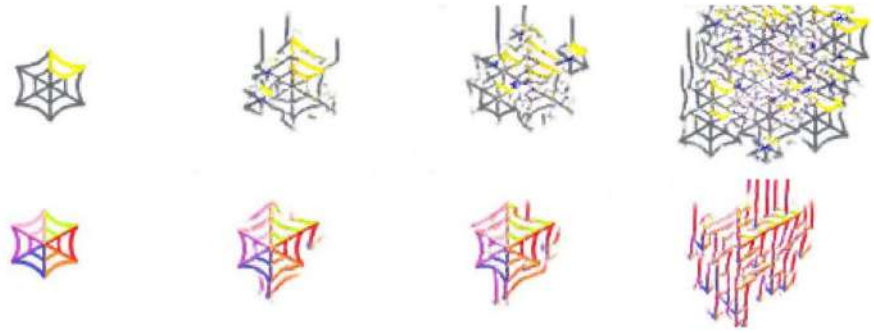

Figure 108: Expansion behaviours for the differently colored spiderweb after 100, 250, 300 and 600 iterations.

The first experiment focused on training the NCA to grow specific target images, following the plain method proposed by [118]. While the initial results did not meet expectations for stable image generation, they provided valuable insights into the NCA's learning process. We trained the NCA on various shapes, including a spiderweb (Fig. 107, 1st column), using a single seed image as the starting point for evolution. The key observation was that after reaching the target image, the NCA consistently continued to expand and distort the figure in subsequent iterations, rather than stabilizing. This uncontrolled expansion behavior varied depending on the specific shape, highlighting how different image characteristics influence the automaton's growth and eventual degradation.

To better understand the behavior of NCAs, experiments were conducted on various shapes [33], which revealed key insights into how these systems grow and self-destruct. In relation to the spiderweb shape (Fig. 107, 1st column), note that: the pattern tends to repeat over time, with new "seeds" emerging from overgrowth, and a strong visual tendency toward diagonal formations, similar to the gliders seen in Conway's Game of Life.

Furthermore, it can be observed that the specific properties of the evolved image significantly influence its growth and alteration. For instance, as shown in Fig. 108, different colorings of the spiderweb lead to distinct emergent patterns. This indicates that the structure of the image, including its internal patterns and color distribution, plays a crucial role in how the NCA learns and evolves.

The behaviors of other shapes, including the *Ticket*, *Alien*, *Sleigh*, and *Caterpillar*, were also extensively analyzed. These figures exhibited unique expansion and self-destruction characteristics, providing further insight into the influence of image properties like symmetry, complexity, and contour on NCA dynamics. Detailed observations and experimental results for these figures are presented in the original study [33].

The second experiment using the pooling strategy yielded the desired results. The NCA trained by this strategy is able to grow from one living cell and, once reaching the targeted image, stabilizes and the image is not expanded further. As for the comparison with the first approach without pooling, Fig. 109 illustrates how the NCA without pooling expands the spiderweb image in a span of 1000 steps post-training and how the NCA

using the pooling algorithm stabilizes, keeping the figure unchanged even while further evolving the NCA.

Figure 109: Comparison of an organism grown after training with Algorithm 9 (top) and respectively the *Algorithm 4 - Pooling Strategy*, from [33] (bottom) steps 0, 50, 100, 150, 250, 350 and 1000.

To further analyze how the pooling training strategy influences NCA behavior, a sample batch (batch 700 out of 8000) was examined during training, as shown in Fig. 110. The first row displays the randomly selected images from the pool before evolution, while the second row shows their state after the NCA applied its update rule.

Figure 110: Batch 700. Based on the approach presented in *Algorithm 4 - Pooling Strategy*, from [33], the randomly selected pool of individuals before and after training of batch number 700.

Observations from this batch confirm the advantages of the pooling approach:

- The neural network successfully guides the evolution of a seed towards an incomplete target image (1st column).
- Importantly, for partially formed images, the neural network learns to stop expanding, maintaining the existing structure (columns 2-7).
- An initial challenge, evident in the last column, was the simultaneous evolution of multiple intertwined spiderwebs from a single input, indicating a duplication problem. However, this issue was observed to be remedied in later stages of training, demonstrating the strategy's effectiveness over time.

The experiments also investigated the influence of filter orientation on NCA behavior. When the standard $Sobel_x$ and $Sobel_y$ filters were replaced with their 45° rotated counterparts, the NCA exhibited a similar overall behavior to the original setup. Figures were still distorted after a number of iterations, with the primary difference being the direction of

these distortions. In both cases, the pooling algorithm *(Algorithm 4 - Pooling Strategy*, from [33]) remained necessary for stabilization.

However, a promising outcome was achieved when the Laplace filter was used in conjunction with the $Sobel_x$ and $Sobel_y$ filters. This combination significantly improved the NCA's performance. This suggests that providing additional, complementary features through the Laplace filter enables the NCA to learn a more reliable evolution path towards the target image.

Finally, this section specifically investigated the behavior of NCAs designed to represent 2D objects, focusing on two core issues: what causes image self-destruction after prolonged evolution and what image characteristics primarily influence its growth. The main advantage of such an automaton is its ability to reproduce an entire shape from a single seed, opening avenues for applications like automated structural repair or pattern reproduction from sparse starting points. However, a current limitation is the need for separate training for each distinct form. A significant contribution of this research is the comprehensive analysis of various image properties, including symmetry, complexity, orientation, and contour, and their impact on the growth principles of the considered images. We observed that growth maintains the original pattern, with colors adhering to their close neighbors rather than spreading randomly. Another important contribution is the detailed analysis of how different filters affect the stabilization process. These findings highlight the essential role of oriented high-pass filters in enabling successful evolution within this NCA architecture, offering valuable insight into how filtering choices directly impact system behavior. The full methodology and experimental results outlined in this section (**I.1**) have been published in *IEEE Access*, a peer-reviewed journal ranked in Q2 according to journal metrics.

II. Signal-Responsive NCAs for versatile multi-texture synthesis

Problem statement

In the field of computer graphics, texture synthesis plays a fundamental role in creating realistic and diverse visual content, being particularly important in industries such as game and film production. While manual texture creation methods are intensive, there is a growing demand for efficient and automated generation techniques. A recent and promising approach in image regeneration and texture synthesis is the use of NCAs [118], [124]. NCAs offer distinct advantages, including relatively reduced computational costs, efficient sampling, and the ability to model complex global patterns from local interactions using a compact internal representation [115], [124], [128]. Unlike many other parametric solutions based on neural networks, which use a large number of parameters (typically $10^5 - 10^6$) for texture generation [63], [171], NCAs maintain a high degree of compactness.

However, a significant limitation of current NCA-based architectures for texture generation lies in their design: each trained NCA is typically specialized to generate a single target texture [124], [117]. While effective for individual instances, this "one-automaton-

per-texture" paradigm inherently becomes demanding and inefficient when aiming for large-scale procedural generation of diverse texture libraries. The necessity to train and manage a distinct automaton for each desired texture severely restricts the scalability and practical utility of NCAs for applications requiring a wide range of textures. This lack of generalization capability across multiple textures within a single model represents a critical bottleneck, preventing the full exploitation of NCAs' potential as a versatile and efficient texture synthesis mechanism. Addressing this shortcoming is crucial to enhance the applicability of NCAs in generating varied and complex visual assets without incurring excessive computational or management overhead.

Methodology and experimental validation of Signal-Responsive NCAs

This section aims to present the foundational concepts and innovations underlying an innovative approach, specifically designed to overcome the inherent limitation of single-texture specialization in current NCA architectures. We begin by outlining the architecture and operational principles of a standard NCA tailored for single-texture generation, with particular emphasis on the loss function and the cell perception mechanism. Following this, it was introduced a novel NCA model, purpose-built for multi-texture generation, which represents the central contribution of this section.

The experimental methodology relies on a proposed genomic encoding scheme embedded within the cells' initial state, allowing for the concurrent development of multiple textures. We provide a detailed justification for the choice of this encoding strategy, alongside an in-depth discussion of the model's architectural design and training procedure. In addition, a unique texture interpolation behavior exhibited by the NCA is briefly presented, along with its implications and potential applications.

Single texture generation

Building on the architecture in [124], this section establishes a baseline for extending NCA to multi-texture generation by exploring optimal perception kernels, network architecture, and loss functions. Each cell's information is encoded in an n_s-dimensional state vector, a model hyperparameter. This vector consists of two main components: the first three values represent color channels, correlating with RGB to assign a pixel's color, while the remaining n_h values constitute hidden channels ($n_h = n_s - 3$) designed to facilitate cell communication. In experiments, common values for n_h were 9, 10, or 12, with specific values detailed per experiment.

For single texture generation, an NCA based on the architecture proposed in [124] was used. An NCA pass, applied at the transition from time step t to $t+1$, consists of two major steps:

- *Perception:* This stage employs four fixed convolutional kernels, namely `Identity` (I_3), `Sobel_x` (S_x), `Sobel_y` (S_y), and a 9-point `Laplacian` (Lap) to convolve with each

channel of the cell's state. These kernels, specifically defined as:

$$S_x = \begin{bmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{bmatrix} \quad S_y = \begin{bmatrix} -1 & -2 & -1 \\ 0 & 0 & 0 \\ 1 & 2 & 1 \end{bmatrix} \quad Lap = \begin{bmatrix} 1 & 2 & 1 \\ 2 & -12 & 2 \\ 1 & 2 & 1 \end{bmatrix}$$

provide the neural network with contextual information from the cell's neighborhood. Circular padding is applied to maintain image size and ensure visual tileability.

- *State processing and update:* The concatenated outputs of the perception kernels are fed into a neural network. This network consists of 1×1 convolutional filters, ensuring that each cell's state update is based solely on its perceived input, simulating independent, per-cell processing. The network generates modification values for each cell's state, which are then applied stochastically (e.g. to half of the cells at each timestamp) to promote self-organization and break symmetry. The architecture allows for training on 128×128 examples and generating textures of arbitrary rectangular dimensions during inference.

  The neural network employed for learning the update rule is displayed in Fig. 111b.

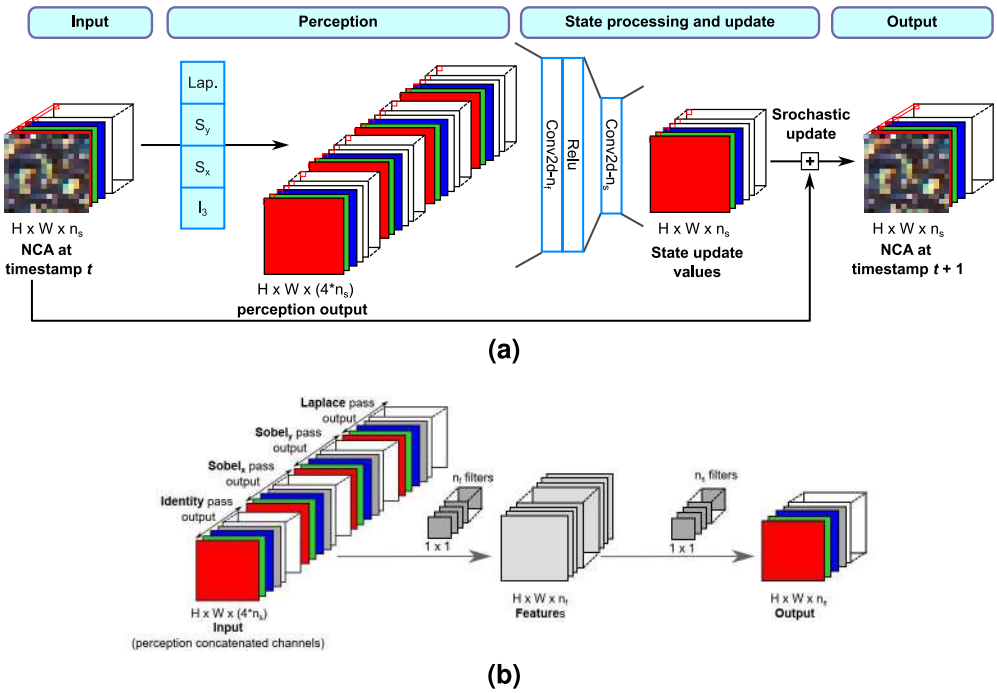

Figure 111: a) One step of the NCA evolution. b) Neural network modeling the update rule from perception output.

The training process involves optimizing the neural network that defines the NCA's update rule. The NCA is initialized with a uniform state and allowed to evolve for a random number of steps (between 32 and 96), as utilized in [124]. The training process is illustrated in Fig. 112. The NCA runs iteratively through the neural network for t steps, as illustrated in Fig. 111a. Each channel of the NCA is convolved with 4 filters, the results

are concatenated and passed through the neural network and the obtained update values are then added to the original NCA at timestamp t, in order to generate the NCA at timestamp $t + 1$.

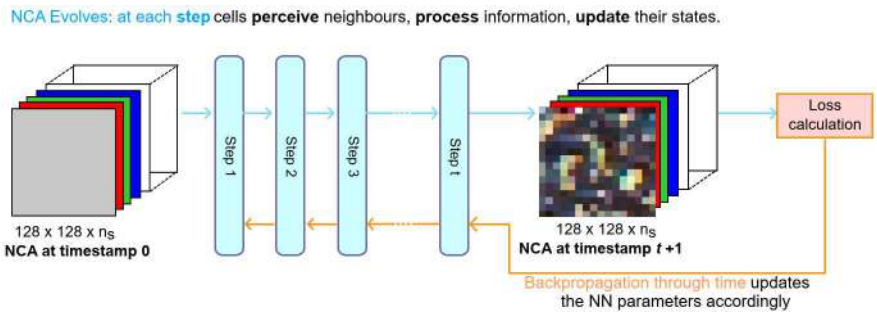

Figure 112: A training step.

The loss is then calculated and the neural network's parameters are updated by backpropagation through time. The quality of the generated texture is then measured against a target example, and backpropagation is used to update the neural network's parameters. The objective is to imitate the style of a given example, not to create a pixel-perfect copy. This is achieved by using an "observer" or "differentiable texture discriminator", specifically a pre-trained VGG16 network [117], [124]. VGG16 is chosen due to its superior performance in style transfer tasks, as it effectively captures style features from selected layers (`conv1-1`, `conv2-1`, `conv3-1`, `conv4-1`, `conv5-1`) [63]. The loss function matches the distributions of these feature maps extracted from both the generated and example images. Fig. 113 illustrates the loss calculation process.

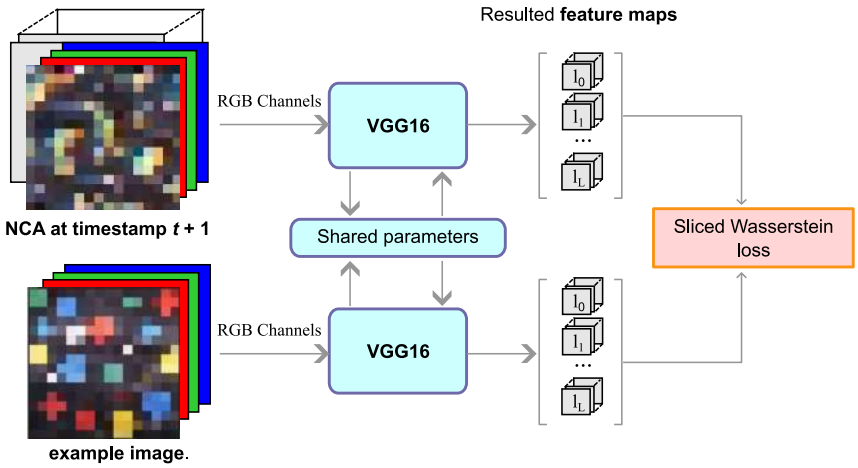

Figure 113: The loss calculation process.

The primary loss used is the Sliced Wasserstein Loss (SWL) [71], which is noted for its effectiveness in capturing complete feature distributions. The calculation details of the total loss are presented in the original paper, see [31]. Furthermore, Optimal Transport Loss (OTT) [115] was explored as an alternative to generate regular patterns, with comparative results discussed in later subsections. The corresponding loss decay curves for both loss functions during training are shown in Fig. 114.

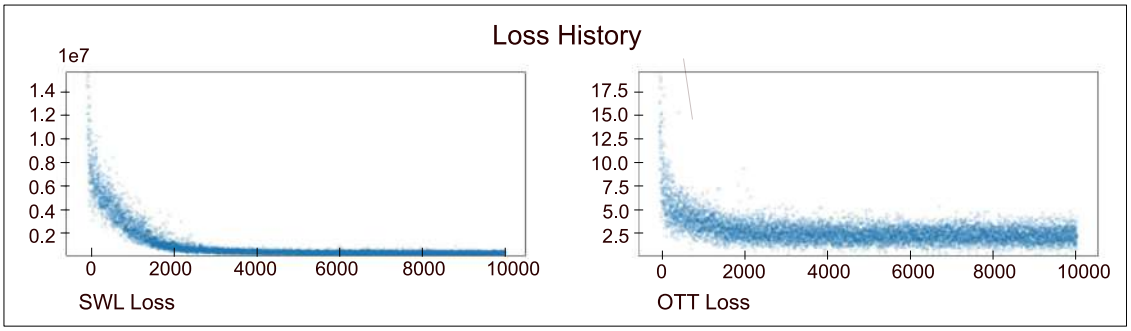

Figure 114: Loss decay in NCA training: SWL (left) vs. OTT (right).

To ensure the long-term stability and consistent behavior of the automaton beyond the initial training steps, a sample pool strategy is employed [118], [124]. A pool of 1024 future textures and seed states is maintained. During training, batches of 8 elements are selected from this pool, and after a training step, the updated states are returned to the pool. To encourage seed evolution and prevent training on irrelevant hallucinations, the element with the highest loss in a batch is replaced with a newly initialized uniform seed.

Multi-texture generation

This section introduces the core concept of multi-texture generation by integrating "internal signals" [163] into the NCA model. This integration is achieved by assigning a special meaning to a subset of the cell's hidden state channels, specifically designating them as "genomic channels" (n_g), distinct from the communication channels (n_c). The experiments explore the use of $n_g = 1$, $n_g = 2$, and $n_g = 3$ genomic channels, allowing a single NCA to generate 2^{n_g} distinct textures. A binary encoding strategy is employed for these genomic channels, where their initial values are set to either 0 or 1. This choice is particularly suitable given the use of an overflow loss that encourages channel values to remain within the $[-1, 1]$ interval. This binary encoding not only facilitates the generation of multiple textures but also inherently supports interpolation between learned textures, a key advantage discussed further in the context of regeneration.

The process of initializing the NCA with a specific genomic signal is essential. As illustrated conceptually in Fig. 115, for an expected texture (e.g., "frilly" or "stratified"), the corresponding binary code is assigned to the genomic channels of every cell at timestamp 0. This initialization is formalized by the `seed_of_genome` function (Algorithm 10), where the last n_g state channels of the `seed` array are set to the binary representation of the target genome index, while all other channels are initialized to zero. In Fig. 115 the tuple for each texture corresponds to the (g_1, g_2, g_3) channels in the seed state. In order to generate one of the given textures at inference, all cells should have the state vector with all values set to 0, except the last three positions, which are set to the values corresponding to the respective genome.

After this initial setup at timestamp 1, the automaton is left to evolve autonomously; no further restrictions or modifications are imposed on any channels, beyond the general

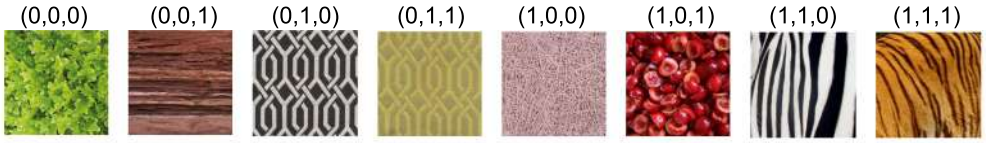

Figure 115: Examples of textures generated by a 3-genome NCA architecture.

Algorithm 10 Function `seed_of_genome`.

Input: h, w - height and width of cell plane (size of texture), n_h - number of cells' state hidden channels, with n_g, genome channels, i_g - expected genome index
Output: One NCA state of the specified genome

$seed \Leftarrow zeros(h, w, 3 + n_h)$
$g \Leftarrow to\_base\_2(i_g, n_g)$ ▷ binary representation of g on n_g bits
$seed[:, :, -n_g :] \Leftarrow g$
return $seed$

encouragement of low values via the overflow loss. The success of this approach hinges on the NCA's ability to "understand" and intrinsically preserve this genomic information throughout its evolution to generate the desired texture.

Beyond generating textures based on discrete example images (2, 4, or 8), this approach investigates the interpolation capabilities of NCAs. Interpolation, in this context, refers to the creation of blended textures, facilitating a smooth visual transition between two distinct learned examples. For instance, if an NCA has learned a "frilly" texture (e.g., genomic code 000) and a "stratified" texture (e.g., 001), interpolation allows the generation of a hybrid texture (e.g., a "mossy wood" texture from a combination of two examples) by setting an intermediate genomic channel value (e.g., $g = (0, 0, 0.5)$).

Technically, interpolation involves creating a new genomic code $g_3 = (c_0, c_1, \ldots, c_{n_g})$ from two learned genomic codes $g_1 = (a_0, a_1, \ldots, a_{n_g})$ and $g_2 = (b_0, b_1, \ldots, b_{n_g})$. Where corresponding bits in g_1 and g_2 are identical ($a_i = b_i$), c_i retains that value ($c_i = a_i$). Where they differ ($a_i \neq b_i$), c_i is set to an intermediate floating-point value (e.g., 0.25, 0.5, 0.75). While this allows for blending between any two learned textures, experiments indicate that interpolations requiring multiple intermediate values in the genome can lead to less stable and predictable results. The specific algorithm for generating an interpolated genome is detailed in *Algorithm 2 - Interpolated genome initialization* from the original paper [31]. This interpolation behavior is a unique feature of the NCA architecture within the domain of cellular automata for texture synthesis, as previous NCA models typically generate only single textures. Note that texture interpolation is performed during the inference phase and does not necessitate any modifications to the NCA's training strategy.

The training strategy for multi-texture generation employs an adapted pooling mechanism to ensure stability and effective learning across all textures. The pool is initialized with an equal distribution of seeds for each target texture. Note that when replacing elements in the pool, the highest-scoring (worst performing) state *belonging to a specific genome* is replaced with its original seed. This replacement cycle iterates through each genome sequentially (e.g., replacing for genome g_0 in one batch, then g_1 in the next, and so on). This method prevents more intricate or harder-to-learn textures, which might con-

sistently yield higher losses, from being perpetually replaced, thereby maintaining their presence in the training process and ensuring their eventual stability during inference.

The detailed pooling-based training strategy is outlined in *Algorithm 3 - Pooling strategy based training for multi-texture generation* from the original paper [31]. This algorithm describes the pool's initialization, where each element tracks both the NCA state and its corresponding genome index (essential as genome channels evolve). The training loop involves randomly selecting a batch, identifying the worst-performing example for the currently considered genome, replacing it with its seed, running the batch through the NCA for updates, calculating the loss, adjusting NCA weights via backpropagation, and finally returning the updated states to the pool. This iterative process allows the algorithm to robustly evolve and refine multiple textures within a shared automaton.

Experimental validation

A series of experiments was conducted to analyze the impact of hyperparameters (first hidden layer size, loss function) and genome influence on texture generation, also evaluating the NCA's texture interpolation capabilities. All images were sourced from the Describable Textures Dataset (DTD) [41] and VisTex Database [114]. Experiments were performed on a T4 GPU, with each lasting up to 2 hours.

Single-texture generation experiments focused on the impact of different perception filters, utilizing a pool of 1024 images, a batch size of 8, and 5000 training epochs. Results indicate that hardcoded $Sobel_x$, $Sobel_y$, and $Laplace$ filters provide sufficient neighbor representation, leading to significantly better texture quality compared to learned kernels, without adding unnecessary complexity. This study's primary contribution, however, is the use of a single NCA for multi-texture generation via genomic channels, alongside exploring novel NCA applications.

The investigation into multi-texture generation and interpolation involved analyzing various NCA architectures, their long-term stability, and interpolation capabilities between learned genomes. All NCAs were trained for 10000 epochs, using a pool size of 1024 samples and a batch size of 8 images. Perception relied on `Sobel` and `Laplacian` filters. The state's hidden channels (n_h) were divided into communication (n_c) and genomic (n_g) channels, and n_f denotes the number of filters in the neural network's first convolutional layer. Table 29 summarizes the performed multi-texture experiments.

Table 29: Multi-texture initial experiments

No. genomes	Experiment	n_h		n_f	Aim
		n_c	n_g		
2	G2Feasible (G2F)	9	1	96	feasibility of model
4	G4Similar (G4Sim)	9	2	128	similar (sim.) textures
	G4Different (G4Diff)	9	2	128	different (diff.) textures
	G4Structured (G4Str)	9	2	128	structured textures
8	G8Large (G8L)	9	3	128	diff. + sim. textures
	G8Medium (G8M)	6	3	70	small model
	G8SmallNoR (G8SNR)	3	3	30	xs model (no regen.)

The experiments covered various aspects of texture similarity and regularity:

- *Experiment G2Feasible (G2F):* Tested the feasibility of genomic coding for differentiating between two textures. Results were satisfactory, confirming the NCA's ability to learn and differentiate.
- *Experiment G4Similar (G4Sim):* Explored generating four similar textures, showing the NCA learns to differentiate but struggles with long-term stability (see Fig. 116a).
- *Experiment G4Different (G4Diff):* Aimed at generating four distinct textures, successfully demonstrating the NCA's independence from shared features. However, it revealed a limitation in achieving global organization for highly structured textures (e.g., grid, see Fig. 116b).
- *Experiment G4Structured (G4Str):* Followed up on the global organization limitation, training with highly structured patterns. This confirmed the NCA's difficulty in achieving global organization, though it could imitate smaller patterns (see Fig. 116c).
- *Experiments G8Large (G8L) and G8Medium (G8M):* Focused on 8-texture generation using diverse artificial and real-life examples. G8M optimized G8L's parameters, achieving similar results with significantly fewer parameters (4270 vs 10k) (see Fig. 116d and Fig. 116e).
- *Experiment G8SmallNoR (G8SNR):* Tested the smallest possible neural network size for 8-texture generation, foregoing regeneration. It resulted in a 1500-parameter NCA stable up to 500 steps, highlighting a trade-off with long-term stability (see Fig. 116f).

Fig. 116 visualizes the textures generated by the proposed multi-texture experiments (*G4Sim*, *G4Diff*, *G4Str*, *G8L*, *G8M*, *G8SNR*), including variable-sized rectangular outputs. The experiments successfully demonstrated the automaton's ability to evolve numerous textures based on internal genomic signals. A key finding was that many training experiments used unnecessarily large neural networks; for instance, the 10k-parameter G8L architecture yielded similar output quality to the 4k-parameter G8M, as seen in Fig. 116. However, this parameter reduction came at the cost of generation speed, with smaller architectures requiring significantly more steps to restore textures (e.g., G8M taking 420 steps compared to G8L's 180). Further reduction to a 3300-parameter architecture also succeeded in generating 8 textures, but with even slower generation (700 steps) and long-term instability, where genomes became corrupted around 1000 steps.

Additionally, in Fig. 116 the small samples in the top row correspond to the training examples for the 8 genomes. Below are the generated textures for each genome provided by these experiments. Two examples of the same texture can be seen in (g), but with different sizes 100×200 pixels and 550×100 pixels (rotated by 90^0). At inference, while cellular automata can run indefinitely, these experiments showed that textures typically maintain stability for up to 6000 iterations (some even longer) before genome corruption and blending with other textures occur.

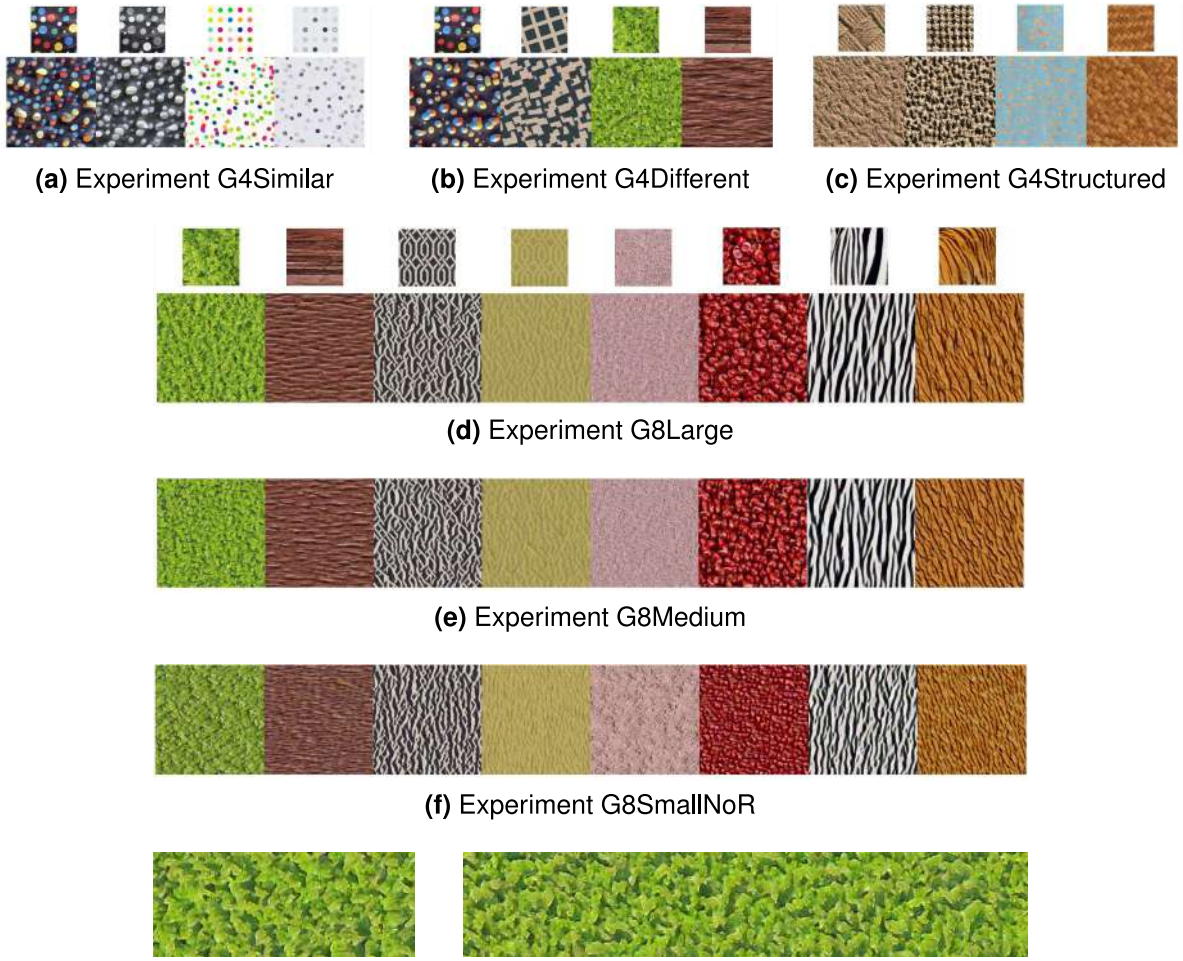

(a) Experiment G4Similar **(b)** Experiment G4Different **(c)** Experiment G4Structured

(d) Experiment G8Large

(e) Experiment G8Medium

(f) Experiment G8SmallNoR

(g) Textures of variable size

Figure 116: The results for multitexture generation, using the architectures as in Table 29, for 4 and respectively 8 different textures.

A key generalization capability of the NCA is its ability to perform texture interpolation, a niche but valuable area in computer graphics. This is unique to the proposed architecture, as other texture synthesis NCAs are limited to single-texture generation.

Fig. 117a illustrates interpolation for the G4Sim experiment, blending a colored and grayscale polka-dot texture. The NCA successfully creates a visually coherent blend that preserves the properties of the original genomes, avoiding artifacts like mixed colors within a single dot. This supports the idea that cells "find an algorithm" for pattern generation [124]. Further examples in Fig. 117b demonstrate the utility of this interpolation for more distinctive textures, opening new avenues for NCA development in texture generation. It is noted that the choice of loss function significantly influences interpolation behavior, as explored in other works [77].

Experiments led to two primary research questions:

1. To what extent does the automaton preserve genomic information within its cells during evolution, and how does this enable the generation and stable differentiation of various textures?

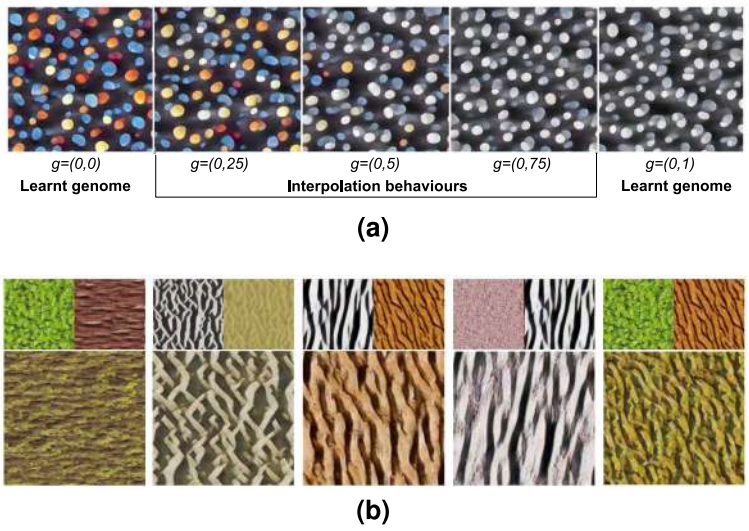

Figure 117: a) Interpolation of learnt genomes for the G4Sim experiment. b) Interpolated textures for the 8-genome experiment G8L, genomes initialized left to right with: $g = (0, 0, 0.5)$, $g = (0, 1, 0.5)$, $g = (1, 1, 0.5)$, $g = (1, 0.5, 0)$ and $g = (0.35, 0.35, 0.35)$

2. What is the influence of the loss function on improving the generation of highly-structured textures that necessitate broad-image communication?

Both of these fundamental questions are thoroughly explored within the original study [31], co-authored by the author of the present thesis. Shortly, related with with the first question, the automaton demonstrably preserves and utilizes genomic information within its cells during evolution, enabling the generation and stable differentiation of various textures. This is achieved through its capacity for active differentiation between similar textures and the suppression of inconsistent features, both guided by the assigned genome. Furthermore, the genomic channel facilitates the feature-specific encoding of texture characteristics, such as assigning lower values for particular visual elements. Ultimately, this encoded information then serves as an internal blueprint to maintain texture fidelity and stability over time. For the second one, the experiments showed that the loss function directly dictates the automaton's ability to capture global structure in highly-structured textures. SWL excels at fine details but struggles with broad communication, leading to a weaker global representation. In contrast, OTT Loss specifically enhances broad communication, making it superior for generating regular, highly-structured patterns by prioritizing global coherence.

Lastly, a quantitative evaluation of the generated textures is provided, complementing visual assessments. The goal is not to reproduce input textures identically but to capture their overall appearance and structural characteristics. Table 30 presents numerical results for NCAs trained with both SWL and OTT losses, evaluated using three key metrics. The *Structural Similarity Index (SSIM)* $\uparrow$ quantifies structural and perceptual similarity between images. The *Learned Perceptual Image Patch Similarity (LPIPS)* $\downarrow$ measures perceptual similarity based on deep neural network features. Finally, the *Gram Matrix Distance (GMD)* $\downarrow$ assesses stylistic similarity, a common measure in style transfer

tasks. These metrics were averaged over 10 independent NCA evolutions for each of 6 regular and 6 irregular textures. Results are reported separately for regular, irregular, and all textures.

The quantitative results support visual observations. Low SSIM values confirm that the NCA does not produce exact replicas of the input textures. LPIPS values are comparable to those reported in the literature [148], [144], indicating successful capture of overall style. GMD values remain within acceptable fidelity limits.

The assessment reveals that OTT loss generally yields better results for regular textures, while SWL performs better on irregular patterns. This suggests that OTT loss compensates for SWL's limitations in accurately reproducing textures with regular, relatively large patterns, making it preferred for such cases. For a broader approach, SWL covers most general examples, and a hybrid approach could offer improved results for both regular and irregular patterns in future work.

Table 30: Quantitative assessment of texture generation by the NCA.

Metric	SWL loss			OTT loss		
	All textures	Irregular textures	Regular textures	All textures	Irregular textures	Regular textures
SSIM ↑	0.0898	0.0917	0.0767	0.1022	0.0735	0.1309
LPIPS ↓	0.4313	0.4472	0.4153	0.4377	0.4985	0.3769
GMD ↓	0.0121	0.0147	0.0095	0.0197	0.0061	0.0334

This chapter focuses on addressing the core issues identified in the original article [31]. This study significantly advances the field of NCA by demonstrating how a single compact model can generate multiple high-quality textures and even smoothly interpolate between them, a capability previously unachievable with traditional NCA approaches. By introducing and effectively utilizing internal genomic signals, it has been shown that these self-organizing systems can learn to differentiate and stably produce diverse patterns based on simple coded instructions. This work not only refines the training methodologies for greater stability and accuracy across various texture types, but also provides compelling evidence that the automaton actively preserves and utilizes this genomic information throughout its evolution, effectively translating abstract codes into complex visual structures. Ultimately, this research positions NCAs as a highly versatile and powerful tool for texture synthesis, offering an extensible and efficient solution for computer graphics and beyond. This article is currently under revision (minor revision, step II) for publication in the Q1 journal, *Scientific Reports*.

(B-ii) The evolution and development plans for career development

This chapter outlines a strategic plan for future academic and professional growth, building on the insights discussed in the preceding sections. It outlines targeted objectives for teaching, research, and institutional engagement, all aimed at ensuring sustained and meaningful contributions to the academic community.

Specifically, this chapter focuses on three closely connected areas that will guide future academic work and involvement:

I. Future research directions - This section explores key thematic areas that require further investigation, aiming to expand current knowledge and address new challenges in the field of career development. Here I address the immediate research topics I plan to explore together with my future doctoral students. While some of these themes have already been introduced at the bachelor's and master's thesis levels, my goal is to advance them to a higher level of academic research, fostering more comprehensive and impactful studies. Special emphasis will be placed on interdisciplinary collaboration, practical applications, and ensuring alignment with evolving societal needs.

II. Didactic evolution - This section focuses on teaching methods and learning approaches to prepare students for a fast-changing job market, where generative AI and emerging technologies are setting the pace and shaping industry demands. I strongly think that the future of teaching lies in adopting student-centered strategies that integrate digital tools, experiential learning, and the development of critical thinking skills.

III. Academic development plans - The final section outlines a personal roadmap for sustained academic growth, focusing on leadership in research and scientific events, curriculum innovation, collaboration and mentoring, as well as active involvement in academic governance and the targeted dissemination of knowledge within scholarly and professional communities.

I. Future research directions

Drawing from my established research foundation and ongoing projects, I am focusing on expanding into strategic areas that combine AI with practical applications. These directions emerge naturally from my current work while pushing into emerging technologies that show significant potential. One main research direction focuses on the use of explainable AI in agricultural contexts, aiming to improve transparency, trust, and decision-making in data-driven crop monitoring and management. The second direction investigates Agentic and Multi-Agent AI systems in complex engineering domains, with an emphasis on autonomy, coordination, and adaptive behavior in real-world industrial settings. Both directions open up a wide range of specific challenges: technical, methodological,

and interdisciplinary, which I plan to explore further through targeted research initiatives and dedicated doctoral supervision.

Explainable AI in Agriculture

This direction represents a natural evolution of my current work in agricultural applications, focusing on developing transparent and interpretable AI solutions. The foundation for this work is already established through:

- Solid foundation in both theoretical and applied machine learning applications and algorithms.
- Active participation in AI4AGRI initiatives and summer schools.
- Current research in agricultural data analysis and modeling.
- Initial exploration of Explainable AI (XAI) methods evolved through the publication of three conference papers in 2025 [174], [175], [162] addressing key aspects of model interpretability and pattern analysis. This progression was further reinforced by supervising two undergraduate theses: one centered on the development of a web-based platform for automated model interpretation, and the other on deep learning techniques for violence detection. The latter integrated Grad-CAM (Gradient-weighted Class Activation Mapping) [195], [150] and introduced an attention-guided cropping mechanism, achieving over 80% across multiple evaluation metrics [175], including with compact models.
- Existing collaborations with domain experts in precision agriculture.
- Transferable expertise in explainable AI methods, developed during my work on deep learning for violence detection.

In recent decades, AI has evolved from simple rule-based systems to complex machine learning models and deep neural networks. While these advanced AI systems achieve impressive performance in tasks like image recognition, natural language processing, and autonomous decision-making, they often operate as "black boxes" providing accurate predictions without explaining their decision-making process [4] [166].

This lack of transparency presents significant challenges, particularly in critical domains like agriculture, healthcare, or finance, where understanding the reasoning behind AI decisions is crucial. Explainable AI (XAI) [47] [2] emerges as a response to these challenges, aiming to develop methods and techniques that make AI systems more transparent and interpretable while maintaining their high performance. Modern XAI methods are structured around the types of questions users might ask, such as: "Why was this decision made?", "Why not another outcome?" or "How can I achieve a different prediction?". In agricultural contexts, additional domain-specific questions emerge: "Which visual features in the satellite imagery most influenced the crop health prediction?", "Why did the model

classify this field as water-stressed when recent rainfall data suggests otherwise?", "What would need to change in this field's data for the model to predict optimal harvest time?", and "How does the model distinguish between natural crop yellowing and disease symptoms?". These explanations can take various forms, from highlighting influential input features to generating counterfactual examples, or constructing locally interpretable approximations (e.g., Local Interpretable Model-agnostic Explanation (LIME) [61], SHapley Additive explanation (SHAP) [103]). Studies demonstrate that regions highlighted by activation maps can effectively be used to suppress background noise or emphasize salient regions, thereby improving model focus and training stability [92]. A diverse set of XAI techniques contributes to model interpretability at different levels. Instance-level methods such as LIME and SHAP provide localized insights by identifying key input features that drive predictions [194] [102], while counterfactual approaches like CEM and DiCE explain how minimal input changes could alter model output [46]. Global techniques, including feature importance scores and surrogate models, offer a broader understanding of the model's decision logic. To more rigorously assess the quality of saliency-based explanations, insertion and deletion tests have been introduced [68], which measure the impact of progressively removing or enhancing salient regions on model confidence. Additionally, layer-wise attribution strategies, such as aggregating Class Activation Maps from multiple convolutional layers using weighted averages, max-pooling, or PCA fusion, have been shown to produce richer and more reliable heatmaps [92].

One of the most influential developments in the XAI domain, particularly for CNNs, has been the advent of visual explanation techniques. Among these, the Class Activation Mapping (CAM) [195] technique, along with its numerous derivatives, such as Grad-CAM [150], Grad-CAM++ [36] or Score-CAM [176], has emerged as foundational tools in understanding and interpreting the decisions of CNNs.

The importance of XAI extends beyond technical considerations, encompassing legal and ethical requirements, such as the "right to explanation" under GDPR. Furthermore, model explainability is crucial for diagnostic and improvement processes. Without proper explanations, neural network developers face significant challenges in detecting overfitting errors, hidden data biases, or cases where models make decisions based on irrelevant features ("shortcut learning"). XAI techniques contribute to identifying these issues, thereby supporting model robustness and ensuring fairness in AI systems. This diagnostic capability is particularly valuable as it enables developers to refine models, eliminate biases, and enhance overall system reliability.

Recent literature demonstrates promising applications of XAI in agriculture. For instance, a comprehensive study [110] explored the integration of explainable AI techniques with precision farming, focusing on crop yield prediction and agricultural decision support. The research successfully demonstrated that combining pre-trained architectures with enhanced optimization algorithms effectively addresses computational complexity challenges common in agricultural data analysis. Furthermore, the study showed that the XL-

Net+SVM model achieved superior performance in crop yield predictions compared to traditional approaches. The integration of interpretability frameworks, specifically SHAP and LIME, provided transparent insights into the decision-making process, making AI recommendations more actionable for farmers. Notably, this approach demonstrated significant improvements in both prediction accuracy and resource management while maintaining interpretability, exemplifying how XAI can bridge the gap between complex AI models and practical agricultural applications.

Another significant contribution to the field [170] addresses the critical challenge of crop diversification through an innovative edge computing-based explainable crop recommendation system called AgroXAI. This system leverages IoT, machine learning, and XAI technologies to suggest suitable crops based on regional weather and soil conditions. What makes this approach particularly valuable is its comprehensive integration of multiple explainability methods (ELI5, LIME, SHAP) to provide both local and global explanations of model decisions. The system goes beyond simple recommendations by incorporating counterfactual explainability methods to suggest alternative crops for different regions, thereby supporting agricultural diversity and adaptation to changing climate conditions.

Addressing the crucial aspect of user experience and adoption, recent research [135] introduces AgriUXE, a collaborative platform specifically designed to bridge the gap between complex AI solutions and farmers' practical needs. This platform focuses on enhancing the explainability of both multimodal data and machine learning predictions, recognizing that transparency is key to building trust and encouraging adoption among small and medium-sized farmers. The approach emphasizes the importance of developing transparent data-driven solutions in collaboration with farm stakeholders, demonstrating how optimized user experience can significantly influence farmers' expectations and acceptance of smart farming technologies.

Building on these foundational studies and identified challenges in agricultural AI applications, my research aims to advance XAI methodologies, particularly in agricultural applications where interpretable decisions are crucial for practical implementation. The integration of XAI in agriculture has shown promising results, from enhanced crop yield predictions to sophisticated recommendation systems and improved user experiences. However, there remain significant opportunities for advancement and innovation in this field. Drawing from the successful implementation of platforms like AgroXAI and AgriUXE, while acknowledging current limitations in model interpretability and user adoption, I have identified several critical directions for future research that combine technical innovation with practical applicability in agricultural settings.

To address these challenges and advance the field of XAI in agriculture, my research will focus on the following key directions:

- Developing novel XAI architectures specifically tailored for agricultural challenges, with particular emphasis on handling complex temporal and spatial data patterns in

crop development.

- Implementing attention-based mechanisms for crop growth monitoring, incorporating multimodal data analysis techniques to capture subtle changes in plant health and development.
- Creating hybrid models combining domain-specific agricultural knowledge with deep learning approaches, enhancing both prediction accuracy and result interpretability.
- Developing frameworks for integrating and interpreting diverse data sources including satellite imagery, sensor data, and weather information, providing comprehensive and explainable agricultural insights.
- Establishing robust validation methodologies specifically designed for agricultural AI models, ensuring reliability and trustworthiness of predictions in real-world farming scenarios. This direction will be strengthened through a potential collaboration with the National Institute of Research and Development for Potato and Sugar Beet Braşov, enabling in-situ validation and real-world testing of the developed models.
- Strengthening collaborative relationships with agricultural research institutions to facilitate real-world testing and validation of developed methods.

Agentic AI for Engineering

Generative AI (GenAI) and large language models (LLMs) [173] [29] can significantly assist teams dealing with heavy workloads and urgent tasks. They do this by automating or accelerating various processes. When used effectively, AI is a powerful tool for boosting productivity and maintaining smooth team operations. Specifically, by integrating AI with engineering tools, simulation users can interact with their modeling and simulation environment using natural, everyday language.

This capability is further advanced by concepts like AI Workflows and Agentic AI [129] [186], two emerging areas with significant potential. AI Workflows are automated sequences designed to handle specific, predictable tasks consistently. Agentic AI, however, takes this a step further. Unlike workflows, agents are autonomous. They can adapt, make decisions, learn from experience, and operate with minimal human oversight. They achieve this by using memory, calling tools, and retrieving knowledge to reason and plan tasks flexibly [189]. While agents offer greater independence, their applications often overlap with workflows.

Building on the concept of individual Agentic AI, a Multi-Agent system involves multiple intelligent agents, such as software, robots, or other autonomous entities, that interact and operate independently. These systems are characterized by their agents' autonomy, social ability, reactivity to environmental changes, and pro-activeness in achieving objectives. Multi-Agent systems offer significant benefits, including scalability (easily adding agents), robustness (resilience to individual agent failures), and flexibility for dynamic

tasks. This makes them ideal for complex problems requiring collaborative autonomy towards shared goals, and they often facilitate interaction through intuitive natural language interfaces.

These foundational concepts, particularly Multi-Agent systems, hold immense research potential, especially within complex domains like engineering, where collaborative intelligence can revolutionize traditional workflows. This leads to a compelling research direction: the development of Multi-Agent systems specifically tailored for systems engineering tasks. Imagine a scenario where each agent specializes in a distinct role – perhaps a system architect, a requirements engineer, or a test lead. This paradigm aims to mirror the collaborative dynamics and specialized expertise found in highly effective human teams. Each agent would be equipped with the necessary capabilities, data, and a deep understanding of project objectives. The ultimate goal is to move beyond AI as mere standalone tools, enabling these AI agents to achieve a state of "team flow" and integrate seamlessly, thereby significantly boosting overall team productivity and workflow.

My interest in these transformative AI concepts, especially AI Workflows and Agentic AI, began early on, as their potential started to appear in initial research. As a first practical step in exploring their real-world use, I contributed to the study [157]. This work specifically used LLMs to automatically generate executable metamorphic relations (EMRs) from software requirements, tackling a significant challenge in automated testing. Our approach showed how LLMs, guided by a few-shot prompting strategy, could handle complex testing activities, much like a human software engineer. Validated through a collaborative industry study, this research highlighted how promising AI workflows are for creating understandable and relevant test artifacts, thus demonstrating the clear benefits of AI in complex engineering tasks. Despite its recent publication, this work has already garnered significant attention, evidenced by its high-quality citations, predominantly from A-ranked and even A*-ranked conferences and journals.

In addition to my research work on this subject, I have supervised several master's theses, many conducted in collaboration with *Siemens Industry Software*, all focused on core topics such as AI workflows, agentic AI, and Multi-Agent systems. These projects have consistently received excellent evaluations from their respective evaluation committees, highlighting both their academic quality and practical relevance. This strong track record in mentorship provides a solid foundation for advancing research at the doctoral level. Below, I outline these projects, emphasizing their individual contributions and identifying potential directions for future exploration, including planned improvements and next research steps to be pursued with PhD students.

- *AmesimGPT - Model manipulation with AI (presented in June, 2025)*

  Project summary: AmesimGPT investigates the capabilities of LLMs to manipulate Simcenter Amesim system simulation models using natural language instructions. The goal is to let users, especially beginners, add, remove, and modify components and global parameters without detailed knowledge of the underlying interface or Python API. The

project evaluates and compares several methodologies: few-shot prompting, retrieval-augmented generation (RAG), router-based planning agents, and hierarchical Multi-Agent systems.

Engineering relevance & AI impact:

- Enhances design and prototyping by enabling rapid, iterative modifications of system models through intuitive natural language queries.
- Facilitates optimization workflows, as engineers can quickly adjust model parameters, swap components, or test alternatives via AI-driven suggestions.
- Automates repetitive and error-prone tasks (e.g., component addition, parameter tuning, connection management), reducing human error and increasing productivity.
- Demonstrates the strengths of LLM-based agents in interpreting ambiguous or high-level engineering requests and converting them into concrete model manipulations.
- Showcases Multi-Agent collaboration, where specialized agents (for components, parameters, submodels, etc.) work under an orchestrator, reflecting real engineering workflows.
- The proposed system features human-in-the-loop integration, allowing it to request additional input from users whenever queries are ambiguous or have multiple valid solutions.

Key challenges:

- Lack of domain-specific knowledge: Even advanced LLMs, trained on broad engineering data, lack the specific mapping between engineering concepts and Amesim's component library.
- Ambiguity in natural language: Even with human-in-the-loop, mapping very high-level descriptions to precise model changes is difficult without extensive context or documentation.
- Complexity collapse in multi-step tasks: LLMs struggle with long, interdependent workflows typical of real-world system engineering (e.g., subsystem or system-level modifications).
- Documentation quality and coverage: RAG methods depend on the granularity and accuracy of available documentation.
- Execution overhead: Multi-Agent or planning approaches increase accuracy but introduce latency and complexity.

Future research and exploration:

  - Domain-specific knowledge fine-tuning: Finetuning LLMs with curated datasets, such as real design sessions, model change histories, and expert clarifications, would help models better understand Amesim-specific terminology, components, and workflows. This targeted training allows LLMs to more accurately map high-level user requests to the correct components or parameters, reducing the need for user intervention. Techniques like QLORA [44] make large-scale, memory-efficient fine-tuning feasible, even for very large models.
  - Fine-tuning for tool use and task decomposition: Effective model manipulation often requires planning and executing multi-step, interdependent actions, similar to how human experts approach complex tasks. Training LLMs on expert demonstration trajectories, logs of tool usage, command sequences, and decision points, helps agents not only understand what needs to be done, but how to do it efficiently within the tool's constraints. Methods such as Direct Multi-Turn Preference Optimization (DMPO) [156] align LLMs with expert workflows by optimizing for preferences over entire action sequences, improving robustness and usability in handling complex, multi-step engineering tasks.

- *Interpreting formulas and code with LLMs (presented in June, 2024)*

  Project summary: The project is an AI-powered internal digital assistant developed in collaboration with *Siemens Industry Software* to streamline employee interactions with company products. It leverages OpenAI-based LLMs to provide intuitive access to internal documentation, code interpretation, and formula analysis. The system integrates multiple operational modes (standard, agent-based) and supports both text and image inputs, using technologies like LangChain, ChromaDB, and Streamlit for backend and frontend implementation.

  Engineering relevance & AI impact:

  - The chatbot enhances engineering workflows by interpreting complex formulas (e.g., fluid dynamics equations) and Java code snippets, enabling faster troubleshooting and knowledge retrieval. AI-driven image analysis (using models like GPT-4 Vision and Claude Haiku) automates the interpretation of technical schematics and mathematical expressions, reducing manual effort.
  - The system automates documentation searches, code explanations, and formula parsing, which are traditionally time-consuming tasks. It integrates with *Siemens Industry Software*' internal APIs (e.g., STAR-CCM+) to provide context-aware responses, aligning with industrial simulation needs.
  - The project demonstrates LLMs' ability to parse and generate Java code (e.g., macros for STAR-CCM+), validated through the *Java Agent* and *Java Agent Graph* modules. It highlights multimodal capabilities (text and image) for technical support, combining OCR (EasyOCR) with vision models to analyze formulas.

Key challenges:

- Traditional vision models (e.g., ResNet-50) performed poorly, achieving only 32–44% similarity scores, compared to GPT-4 Vision's 89–93%.
- Difficulty in code contextualization for domain-specific Java APIs.
- High response times with models like GPT-4 Vision, averaging around 85 seconds per request.

Future research and exploration:

- Exploring multiple hybrid models integration using frameworks like RAGAS to evaluate retrieval-augmented generation for code/formula tasks, and adapting AutoGen for Multi-Agent collaboration in engineering automation to enable dynamic task decomposition (e.g., code generation and validation).

- *reSearchGPT: Chat with Large Document Databases using Retrieval-Augmented Generation and LLM-based Topic Modeling (presented in June, 2024)*

  Project summary: This thesis details the implementation of an advanced chat system engineered to improve user interactions with large PDF Knowledge Bases by integrating a Retrieval-Augmented Generation framework enriched with Topic Identification. Advanced RAG leverages Hierarchical Navigable Small Worlds (HNSW) indexing and optimizes search efficiency by strategically reducing the number of database partitions considered during retrieval in response to a user question. To facilitate this retrieval mechanism, the system employs a clustering step that organizes the vector database into topic-specific partitions. The supporting clustering step is designed to balance computational feasibility and accuracy, combining the rapid convergence of KMeans with the detailed segmentation capability of Hierarchical Clustering. Subsequently, LLMs are used to generate descriptive labels for each partition, enhancing topic interpretability.

  Engineering relevance & AI impact:

  - The project contributes to automating engineering workflows by creating a self-organizing knowledge management system that automatically partitions extensive document databases into topic-specific clusters, generates descriptive labels using LLMs, and enables efficient information retrieval with minimal latency. This automation eliminates the need for manual document categorization and search, significantly streamlining research and information access processes. This helps researches explore existing trends, review previous investigations and validate ideas, all within a conversation-like environment.
  - The project demonstrates the capability of GenAI to comprehend and respond to complex queries by retrieving contextually relevant information from large-scale knowledge bases with high precision and minimal latency. It showcases how LLMs can be effectively combined with retrieval systems to generate coherent responses

based on retrieved contexts, while also exhibiting capabilities in topic modeling and automatic generation of interpretable cluster labels.

Key challenges:

- Optimizing the trade-off between context relevancy and retrieval latency for semantic similarity search tasks within large vector databases (example database has 2 million embeddings, summing to ≈23 GB of data).
- Developing a clustering methodology that balances computational efficiency with accuracy, while operating within the memory constraints imposed by a system with 64 GB of RAM.
- Creating an effective partitioning strategy for topic-specific organization of the database.

Future research and exploration:

- Exploration of alternative embedding models to enhance semantic understanding of technical documents, potentially incorporating recent advancements like BERT-based [25] or domain-specific embeddings [158] which have shown superior performance in semantic similarity tasks.
- Refinement of the clustering methodology through parameter optimization and exploration of alternative document clustering algorithms tailored for information retrieval.
- Integration of document metadata (authorship, publication date, title) into the retrieval process to provide additional context and improve result relevance, following approaches from hybrid retrieval systems in information retrieval literature [143].
- Investigation of more complex topic modeling techniques such as BERTopic [67] that combine embeddings with probabilistic topic modeling for enhanced partitioning accuracy.

- *Fine tuning LLMs on large codebases (presented in June, 2024)*

  Project summary: The purpose of the project is training a LLM on an existing project with a large codebase that exceeds GBs of raw text, in order to create new segments of code relevant to the project, explain already existing parts of the project and guide developers through the codebase. So, the newly created model would help developers boost their productivity. The project research multiple approaches for fine-tuning, in order to improve results, such as fine-tuning in a loop, split the fine-tuning in multiples steps such as "prefinetuning" and "finetuning" and so on. The project used Transformers library provided by HuggingFace for loading models and initial fine-tuning processes and the Unsloth library to decrease memory usage for fine-tuning process. For the creation of the dataset Llama-index library for RAG was used. These technologies are a starting point for fine-tuning and training LLMs, supports a variety of training approaches, such

as supervised fine-tuning training or preference optimization fine-tuning and also for building agentic models.

Engineering relevance & AI impact:

- The LLM optimizes development of large-scale projects, by increasing productivity of existing developers of the project, can disseminate knowledge of the project across developers and help with reducing the time of new developers to become productive on the project.
- The project aims to follow a generic process which could be reproduced on any existing large codebase, to create a dataset and fine-tune a pre-trained LLMs, so the newly created model would generate relevant code segments, explain pieces of code, and guide a software engineer through the codebase. Such an LLM is also useful when a software engineer needs to create repetitive pieces of code.
- The newly created LLM could generate new pieces of code, relevant to the project structure and using project specific classes, functions and patterns, document and explain existing code of a big real world application, in order to smooth learning curve or new colleagues who started to work the codebase and also boost productivity of already existing colleagues.

Key challenges:

- Generating a high-quality dataset consisting of question–answer pairs using an LLM, based on project-specific sources like code, documentation, and version control history.
- Designing a resource- and cost-efficient fine-tuning process that minimizes both hardware usage and training time.

Future research and exploration:

Collect and use user preferences in the fine-tuning process to improve the quality of the newly fine-tuned model. And use agentic AI to help create better datasets for training and also for better question answering. Another direction which can be explored is to use different context augmentation methods such as RAG or Cached Augmented Generation (CAG) to improve question answering capabilities. A few relevant articles that serve as foundational references include [137], [6].

Based on these thesis projects, I'm convinced that the intersection of GenAI, Agentic AI, and Multi-Agent systems with complex engineering domains is a highly promising area for advanced research [42] [196]. The "Future research and exploration" sections outlined in each case are not just minor extensions, they represent significant, multifaceted challenges that demand sustained, in-depth exploration. Given their interdisciplinary nature, connecting fields like AI, software and systems engineering, human-computer interaction, and even cognitive science, I see each of these directions as an excellent subject

for doctoral-level investigation. I intend to continue growing this research line with PhD students, focusing on meaningful problems at the intersection of AI and engineering.

II. Didactic evolution

My didactic development is guided by a continuous commitment to quality, adaptability, and relevance in higher education. Since 2008, I have been involved in academic teaching, initially taking on complex seminars and laboratory sessions that were often avoided due to their challenging nature, such as Operational Research, Computational Logic, and Compiler Techniques. Over the years, I have consistently updated both content and teaching methods, trying to emphasize the practical value of each theoretical concept and to engage students from various specializations, including computer science, mechanical engineering, economic informatics, and information technology. I am explicitly committed to ensuring that what students learn in my courses is practical, immediately applicable, and equips them with skills and knowledge they can use effectively in real-world situations and professional environments

Over time, I have consistently adapted my teaching to create more engaging and interactive learning environments. This evolution is reflected in the range of diverse courses I currently teach across various academic levels: for Bachelor's students, these include *Fundamental Algorithms* (1st Year), *Modern C++ - Managing Networking Projects* (2nd Year), *Modern C++ Applied in Artificial Intelligence* (2nd Year), *Object-Oriented Programming in Java* (2nd Year), and *Formal Languages and Compilers* (2nd Year); at the Master's level (1st Year), I teach *Programming in Python for Engineering Applications* and *Data Mining and Data Warehousing*.

For each course, materials are regularly updated and made accessible via the university's e-learning platform, complemented by practical laboratory work designed to reinforce theoretical concepts. A significant aspect of my didactic approach involves collaborating with young teaching assistants, many of whom are former students now pursuing doctoral studies or working in software companies in Brașov. This pedagogical approach has consistently received positive feedback from students, as reflected in high course evaluation scores and direct testimonials, underscoring the engaging and practical nature of the learning experience. Combining academic and industry experience helps create a richer learning environment and provides students meaningful insights into real-world applications. This comprehensive approach has consistently yielded positive outcomes, with many of our graduates securing positions at leading technology companies or successfully pursuing advanced studies at prestigious universities both nationally and internationally.

Since 2023, I am the coordinator of the *Applied Informatics* undergraduate program at the *Faculty of Mathematics and Computer Science*, *Transilvania University of Braşov*. I have been actively involved in modernizing the curriculum by introducing forward-looking subjects in collaboration with both academic colleagues and industry professionals. A

series of recently added courses explore areas like generative AI, large language models, agent-based systems, and intelligent image processing, ensuring students are exposed to the most relevant trends and skills required by the industry. These updates aim to prepare students for cutting-edge fields and ensure strong alignment with industry needs. The detailed curriculum is publicly available https://mateinfo.unitbv.ro/ro/programe-de-studii/121-licenta/625-fisedisciplina-ia.html. Staying in close contact with students from this program and regularly gathering their feedback is essential to keeping the curriculum relevant and effective.

In the future, I intend to focus on improving my teaching skills, encouraging interdisciplinary learning, creating more collaborative classroom environments, and staying actively involved in outreach activities. To support these goals, my didactic development will follow several key directions:

- I propose participating in teaching stages of at least two weeks at international universities to gain exposure to diverse pedagogical best practices and global academic perspectives.
- I aim to further diversify interactive teaching methods, emphasizing creative collaboration and educational partnerships to foster student-centered learning.
- Establishing collaborative relationships with students from other universities to promote broader academic exchange and project-based learning.
- Annual review and update of course syllabus and calendars to ensure the continuous relevance and dynamism of the educational content.
- Publishing all taught courses with ISBN in recognized publishing houses to ensure wider dissemination and academic rigor of the educational resources.
- Continuously involving students in lectures and seminars through methods centered on discovery learning, team-based learning, and group learning. I promote creativity, inquiry, and problem-solving, using digital tools and methods tailored to today's learning environments.

Additionally, building on my research experience in AI applied in precision agriculture (via the AI4AGRI project, with partners from France and Italy), and in machine learning for industrial simulation (through a 9-year collaboration with *Siemens Industry Software*), I intend to develop a new interdisciplinary PhD-level course. This course will cover topics at the intersection of AI and agriculture or engineering, offering an excellent opportunity for doctoral students from different fields to work together on innovative and research-worthy projects.

Beyond classroom teaching, I continuously coordinate over 15 bachelor's and master's theses annually on cutting-edge topics, with many of these projects leading to student publications and positive career outcomes. Moving forward, I intend to progressively raise

the complexity and encourage original research contributions. I also plan to develop a collaborative research team of students and colleagues focused on areas such as extracting meaningful data from multi- and hyperspectral images, applications of cellular neural automata, using AI/ML in advanced engineering problems, and quantum technologies and their applications.

Building on my current experience in outreach and student engagement, I aim to continue and advance my involvement in promoting academic programs in high schools across Brașov and beyond. I want to get students more involved through both academic and extracurricular activities, drawing on my experience with successful initiatives such as hackathons - Transilvania CodeStorm@MI, student research presentation sessions, challenges and events like AFCO. Furthermore, I plan to expand these efforts by fostering more informal spaces for students to share their ideas, ambitions, and career goals, thereby supporting their academic and professional development.

III. Academic development plans

To ensure continuous academic growth and meaningful contributions to my field, I have outlined a comprehensive plan focusing on leadership in research and scientific events, curriculum innovation, collaboration and mentoring, as well as active involvement in academic governance and dissemination of knowledge. The following sections detail the key initiatives and objectives that will guide my professional development in the coming years.

1. Enhancing scientific dissemination and community engagement

   Sharing research findings and contributing to the scientific community are fundamental to my academic development. To this end, I am committing to:

   - Publishing a minimum of two articles annually in high-impact (WoS indexed with high IF/AIS scores) scientific journals in Computer Science and interdisciplinary topics. My strategy will prioritize geographical and thematic diversity of journals, while also supporting Romanian publications.
   - Actively participating in and leading the organization of national and international scientific conferences and events. For 2025, I am involved in two significant events:
     - FES 2025 Conference, where I serve as Technical Program Chair and Track Chair for the "AI and Advanced Analytics in Energy Systems" session.
     - ACM Celebration of Women in Computing (Brașov, Romania, September 17-19, 2025), where I'm a member of the organizing team (Scholarship chair).
   - As a member of the editorial board for the *Bulletin of the Transilvania University of Brașov. Series III: Mathematics and Computer Science* (a fully Scopus

indexed journal), I am personally dedicated to supervising the rigorous review process for computer science papers, ensuring that only high-quality research is accepted. Regular meetings with colleagues ensure continuous progress in enhancing the journal's prestige and scholarly contribution.

- Proactively disseminating research findings through various channels, including academic publications, conference presentations, and public lectures, to maximize their impact and contribute to the broader scientific discourse. My interest extends particularly to channels that also reach students and the general public, fostering scientific literacy and inspiring future generations. As an example of this commitment, this year I delivered a TEDx talk on artificial intelligence, aiming to make scientific concepts accessible to a wider audience. The talk garnered over 300 views and received positive feedback for its clarity and accessibility, demonstrating effective public engagement.

2. Advancing research leadership and innovation

   My commitment to academic growth is primarily rooted in leading and expanding impactful research. This involves:

   - Building on my experience gained from involvement in 12 national/international grants, I actively seek external funding by preparing and submitting research proposals to both national and international calls. In the current year, I have contributed to the submission of 5 additional project proposals: 2 in my capacity as a Siemens representative and 3 as a university affiliate. These initiatives include both continuations of existing or recently completed projects and explorations into new strategic directions. For instance, I have contributed to proposals investigating the application of AI in entrepreneurial contexts, specifically the HORIZON-MSCA-SE-2025 project named *AI-PASS*. *AI-PASS* aims to advance the state-of-the-art by developing and validating an AI-driven adaptive loyalty ecosystem. For this project, we collaborate with partners such as ISCAP-Porto and UA Alicante. Our primary objective at *Transilvania University of Braşov* (where I am part of the team) is to contribute expertise in AI engineering, NLP, multimodal sentiment analysis, and ML applied to marketing. Another key initiative is a EUREKA project, *AI-Driven Marketing and Human Behavior Research in Sports and Health*. This project involves an interdisciplinary collaboration with U.N.E.F.S Bucharest. For this project, I coordinate the UniTBv team responsible for developing multimodal chatbots capable of processing video, images, and text, with integration functionalities for external systems. This contribution will significantly strengthen UniTBv's innovation profile in advanced artificial intelligence technologies.
   - I plan to take on leadership roles within existing research teams and to actively build new collaborative groups that encourage shared inquiry and innovation.

This includes working more closely with my current collaborators and starting new interdisciplinary projects together with researchers from other fields.

- A key objective is to initiate and contribute to Centers of Excellence that address critical societal challenges, thereby establishing impactful research hubs. Demonstrating this commitment, as a coordinator for UniTBv, I contribute to the submission of the *PN-IV-P6-6.1-CoEx2024-0093 project, Excellence research center for e-mobility technologies (T4EMOB)*. This project, focusing on Climate, Energy, and Mobility, in collaboration with the *Technical University of Cluj-Napoca*, was evaluated with a notable 81.5 points. This score placed it among the 36 projects that successfully passed the evaluation threshold out of a total of 109 proposals submitted to UEFISCDI. This experience motivates me to continue developing and leading impactful initiatives in the future.
- A major goal of mine is to start and lead a regional Mathematics & Informatics cluster, with Brașov as its main hub. This cluster will aim to bring together academic, research, and industry efforts across the region. The idea is to boost innovation, talent development, and economic growth in the IT and applied mathematics fields. To get this project moving, I have already initiated collaborative discussions with universities from Alba-Iulia, Târgu Mureş, and Sibiu to build the initial partnerships. Following these, I intend to present the concept to ADR Centru, seeking their strategic support for the cluster's development.
- I continuously seek and pursue novel interdisciplinary research avenues. My expertise in areas such as machine learning for industrial simulation (developed through a 9-year collaboration with *Siemens Industry Software*) and AI in precision agriculture (from the AI4AGRI project) enables me to effectively bridge traditional disciplinary boundaries. My publication record further attests to this commitment, with a significant portion of my articles being co-authored with collaborators from *Siemens Industry Software* and various interdisciplinary research projects.
- Strengthening existing and forging new collaborations with companies to ensure research is relevant, impactful, and translates into practical applications, fostering a good relationship between academia and industry.

3. Contributions to university and curricular advancement

   Beyond research, my academic development plan encompasses active contributions to the growth and strategic direction of the university and faculty:

   - Active contribution to academic governance, as a member of both the Department Council and the Faculty Council, advocating for student needs and contributing to strategic institutional decision-making processes.

- I am committed to significant administrative duties as Vice-Dean overseeing student affairs, internationalization, and quality assurance. In this role, I work to represent student interests in key institutional decisions. A main focus is to increase student involvement by supporting their initiatives, encouraging their active participation in academic life, and creating an inclusive and collaborative environment for everyone.
- Promoting the development of new and innovative study programs, like a Master's program in the department that focused on applying computer science to engineering, aimed at meeting the changing needs of industry and broadening educational opportunities.
- My commitment to university and curricular advancement involves promoting the practical application of AI for daily tasks among faculty and students, extending these technologies beyond traditional IT domains. As a proactive step, I will be teaching a master's course titled "Applications of AI in education and learning" at the Faculty of Psychology and Educational Sciences from Transilvania University of Braşov, starting in the 2025 academic year. This initiative aims to promote collaborative educational development, particularly by exploring the interaction of AI and cognitive abilities, addressing the critical question: *How does AI redefine the essence of cognitive engagement in problem-solving?* within an AI-augmented learning environment.
- Organizing new summer schools, attracting participants from both Romania and abroad, focusing on key areas such as applied mathematics, optimization, and machine learning, to foster specialized knowledge and networking opportunities for students and researchers. This year, I was invited as a speaker at one of the summer schools within the SEEN Project, and I also participated as both a speaker and a member of the organizing team for the AI4AGRI Summer School, contributing to the successful planning and execution of the event.
- Having consistently supervised at least 10 Bachelor's and Master's theses annually since 2018, I am deeply committed to continuing this meaningful contribution. Through guiding students to explore current and challenging research topics, I actively support the development of future researchers and continuously strengthen the faculty's research activities

I will align my professional development plan with the development strategy of the *Department of Mathematics and Computer Science* from *Transilvania University of Braşov*, with key objectives being the maintenance and enhancement of academic and professional excellence, and, not least, fostering a positive and productive collaboration with all my colleagues.

(B-iii) Bibliography

[1] M. Z. Abd Elrehim, M. A. Eid, and M. G. Sayed. “Structural Optimization of Concrete Arch Bridges Using Genetic Algorithms”. In: *Ain Shams Engineering Journal* 10.3 (2019), pp. 507–516. DOI: 10.1016/j.asej.2019.01.005.

[2] A. Adadi and M. Berrada. “Peeking Inside the Black-Box: A Survey on Explainable Artificial Intelligence (XAI)”. In: *IEEE Access* 6 (2018), pp. 52138–52160. DOI: 10.1109/ACCESS.2018.2870052.

[3] T. Adão et al. “Hyperspectral Imaging: A Review on UAV-Based Sensors, Data Processing and Applications for Agriculture and Forestry”. In: *Remote Sensing* 9.11 (2017), p. 1110. DOI: 10.3390/rs9111110.

[4] J. Adebayo et al. “Sanity Checks for Saliency Maps”. In: *arXiv preprint* (2018). arXiv: 1810.03292 [cs.LG]. URL: https://arxiv.org/abs/1810.03292.

[5] Agenzia Spaziale Italiana (ASI). *PRISMA Products Specification Document Issue 2.3*. Technical Report. Rome, Italy, 2020. URL: https://prisma.asi.it/missionselect/docs/PRISMA%20Product%20Specifications_Is2_3.pdf.

[6] AI@Meta. *Llama 3 Model Family*. 2024. URL: https://ai.meta.com/llama/.

[7] G. Aksu, C. O. Güzeller, and M. T. Eser. “The Effect of the Normalization Method Used in Different Sample Sizes on the Success of Artificial Neural Network Model”. In: *International Journal of Assessment Tools in Education* 6.2 (2019), pp. 185–201. DOI: 10.21449/ijate.555666. URL: https://doi.org/10.21449/ijate.555666.

[8] J. Aschbacher. “ESA’s Earth Observation Strategy and Copernicus”. In: *Satellite Earth Observations and Their Impact on Society and Policy*. Singapore: Springer, 2017, pp. 81–86. ISBN: 978-981-10-3713-9. DOI: 10.1007/978-981-10-3713-9_5.

[9] C. Aybar et al. “CloudSEN12, a Global Dataset for Semantic Understanding of Cloud and Cloud Shadow in Sentinel-2”. In: *Scientific Data* 9 (2022). DOI: 10.1038/s41597-022-01878-2.

[10] A. Băicoianu, M. Demeter, and A. Vasilescu. “Innovative Air Quality System with Emergency Notifications”. In: *Bulletin of the Transilvania University of Braşov, Series III: Mathematics, Informatics, Physics* 12.2 (2019), pp. 443–456. DOI: 10.31926/but.mif.2019.12.61.2.22.

[11] A. Băicoianu, R. Ivan, and I. Popa. “Exploiting Cellular Automata and Linear Regression to Predict Disease Spread”. In: *Bulletin of the Transilvania University of Braşov. Series III: Mathematics and Computer Science* 5 (67).1 (2025), pp. 241–258. DOI: 10.31926/but.mif.2025.5.67.1.18. URL: https://webbut.unitbv.ro/index.php/Series_III/article/view/8774.

[12] A. Băicoianu and A. Mathe. “Diagnose Bearing Failures With Machine Learning Models”. In: *2021 International Conference on Innovations in Intelligent Systems and Applications (INISTA)*. Kocaeli, Turkey: IEEE, 2021, pp. 1–6. DOI: 10.1109/INISTA52262.2021.9548506.

[13] A. Băicoianu, C. M. Păcurar, and M. Păun. “A Concretization of an Approximation Method for Non-Affine Fractal Interpolation Functions”. In: *Mathematics* 9.7 (2021). ISSN: 2227-7390. DOI: 10.3390/math9070767. URL: https://www.mdpi.com/2227-7390/9/7/767.

[14] A. Băicoianu and I. V. Scheianu. “Managing and Optimizing Big Data Workloads for On-Demand User Centric Reports”. In: *Big Data and Cognitive Computing* 7.2 (2023). DOI: 10.3390/bdcc7020078. URL: https://www.mdpi.com/2504-2289/7/2/78.

[15] A. Băicoianu et al. “A Machine Learning Proposal for Condition Monitoring of Vehicle Suspension”. In: *2022 International Conference on Innovations in Intelligent Systems and Applications (INISTA)*. Biograd na Moru, Croatia: IEEE, 2022, pp. 1–5. DOI: 10.1109/INISTA55318.2022.9894160.

[16] A. Băicoianu et al. “Structural Optimization Using Genetic Algorithms”. In: *2022 International Conference on Innovations in Intelligent Systems and Applications (INISTA)*. IEEE, 2022, pp. 1–5. DOI: 10.1109/INISTA55318.2022.9894243.

[17] A. Băicoianu et al. “Fractal Interpolation in the Context of Prediction Accuracy Optimization”. In: *Engineering Applications of Artificial Intelligence* 133 (2024), p. 108380. DOI: 10.1016/j.engappai.2024.108380.

[18] A. Băicoianu et al. “Integrated Digital Platform for Marine Energy Management”. In: *2024 IEEE International Conference on Electrical Systems for Aircraft, Railway, Ship Propulsion and Road Vehicles & International Transportation Electrification Conference (ESARS-ITEC)*. IEEE, 2024, pp. 1–5. DOI: 10.1109/ESARS-ITEC60450.2024.10819756.

[19] A. Băicoianu et al. “A Practical Approach to Bearing Fault Detection using Pretrained ResNet18”. In: *Proceedings of the Future Energy Solutions Conference*. Accepted for publication, in press. IEEE, 2025.

[20] A. Băicoianu et al. “DACIA5: A Sentinel-1 and Sentinel-2 Dataset for Agricultural Crop Identification Applications”. In: *Big Earth Data* 0.0 (2025), pp. 1–32. DOI: 10.1080/20964471.2025.2512685.

[21] A. Băicoianu et al. “Innovative Transfer: Harnessing Automotive Technologies for Maritime Progress”. In: *CONAT 2024 International Congress of Automotive and Transport Engineering*. Cham, Switzerland: Springer, 2025, pp. 94–101. ISBN: 978-3-031-77635-9.

[22] K. Bandara et al. “Improving the Accuracy of Global Forecasting Models Using Time Series Data Augmentation”. In: *Pattern Recognition* 120 (Dec. 2021). DOI: 10.1016/j.patcog.2021.108148.

[23] A. Bannari et al. “A Review of Vegetation Indices”. In: *Remote Sensing Reviews* 13.1-2 (1995), pp. 95–120. ISSN: 0275-7257.

[24] E. Bélisle et al. “Evaluation of Machine Learning Interpolation Techniques for Prediction of Physical Properties”. In: *Computational Materials Science* 98 (2015), pp. 170–177. DOI: 10.1016/j.commatsci.2014.10.032.

[25] I. Beltagy, K. Lo, and A. Cohan. “SciBERT: Pretrained Language Model for Scientific Text”. In: *Proceedings of EMNLP*. 2019. arXiv: 1903.10676 [cs.CL].

[26] K. Bhattacharjee et al. “A Survey of Cellular Automata: Types, Dynamics, Non-Uniformity and Applications”. In: *Natural Computing* 19 (2020), pp. 433–461. DOI: 10.1007/s11047-019-09748-4.

[27] R. Bocu, A. Băicoianu, and A. Kerestely. “An Extended Survey Concerning the Significance of Artificial Intelligence and Machine Learning Techniques for Bug Triage and Management”. In: *IEEE Access* 11 (2023), pp. 123924–123937. DOI: 10.1109/ACCESS.2023.3329732.

[28] B. Breckling, G. Pe’er, and Y. Matsinos. “Cellular Automata in Ecological Modelling”. In: *Modelling Complex Ecological Dynamics*. Springer, 2011, pp. 105–117. ISBN: 978-3-642-05028-2. DOI: 10.1007/978-3-642-05029-9_8.

[29] T. B. Brown et al. “Language Models are Few-Shot Learners”. In: *arXiv preprint* (2020). Preprint. arXiv: 2005.14165 [cs.CL]. URL: https://arxiv.org/abs/2005.14165.

[30] D. Căliman, V. David, and A. Băicoianu. “An Analysis of Improving Bug Fixing in Software Development”. In: *Proceedings of the 18th International Conference on Software Technologies (ICSOFT 2023)*. SCITEPRESS, 2023, pp. 470–477. ISBN: 978-989-758-665-1. DOI: 10.5220/0012119500003538.

[31] M.-M. Catrina, I. C. Plajer, and A. Băicoianu. “Multi-Texture Synthesis through Signal Responsive Neural Cellular Automata”. In: *arXiv preprint* (2024). Preprint. arXiv: 2407.05991 [cs.NE]. URL: https://arxiv.org/abs/2407.05991.

[32] S. Catrina and A. Băicoianu. “Quantum Tunneling: From Theory to Error-Mitigated Quantum Simulation”. In: *Advanced Quantum Technologies* 8.1 (2025), p. 2400163. DOI: https://doi.org/10.1002/qute.202400163. eprint: https://advanced.onlinelibrary.wiley.com/doi/pdf/10.1002/qute.202400163. URL: https://advanced.onlinelibrary.wiley.com/doi/abs/10.1002/qute.202400163.

[33] S. Catrina et al. “Learning About Growing Neural Cellular Automata”. In: *IEEE Access* 12 (2024), pp. 45740–45751. DOI: 10.1109/ACCESS.2024.3382541.

[34] M. Čech, A.-J. Beltman, and K. Ozols. “Digital Twins and AI in Smart Motion Control Applications”. In: *2022 IEEE 27th International Conference on Emerging Technologies and Factory Automation (ETFA)*. Stuttgart, Germany: IEEE, Sept. 2022, pp. 1–7. DOI: 10.1109/ETFA52439.2022.9921533.

[35] X. Chai et al. “Deep Learning for Irregularly and Regularly Missing 3-D Data Reconstruction”. In: *IEEE Transactions on Geoscience and Remote Sensing* 59 (2021), pp. 6244–6265. URL: https://api.semanticscholar.org/CorpusID:261340870.

[36] A. Chattopadhay et al. “Grad-CAM++: Generalized Gradient-Based Visual Explanations for Deep Convolutional Networks”. In: *Proceedings of the IEEE Winter Conference on Applications of Computer Vision (WACV)*. Lake Tahoe, NV, USA: IEEE, Mar. 2018, pp. 839–847. DOI: 10.1109/WACV.2018.00097.

[37] G. Cheng, J. Han, and X. Lu. “Remote Sensing Image Scene Classification: Benchmark and State of the Art”. In: *Proceedings of the IEEE* 105.10 (2017), pp. 1865–1883. DOI: 10.1109/JPROC.2017.2675998.

[38] A. Cherdivar, A. Băicoianu, and I. Popa. “Parametric Structural Optimization Using Genetic Algorithms”. In: *Proceedings of the International Conferences on Applied Computing*. IADIS, 2024, pp. 63–74. ISBN: 978-989870462-7.

[39] S. W. Choi et al. “Machine Learning-Based Prediction of Korean Triage and Acuity Scale Level in Emergency Department Patients”. In: *Journal of Physics: Conference Series* 25.4 (2019), pp. 305–312. DOI: 10.4258/hir.2019.25.4.305.

[40] E. Christophe, D. Leger, and C. Mailhes. “Comparison and Evaluation of Quality Criteria for Hyperspectral Imagery”. In: *SPIE Conference Proceedings*. Vol. 5668. Bellingham, WA, USA: SPIE, 2004. DOI: 10.1117/12.587107.

[41] M. Cimpoi et al. “Describing Textures in the Wild”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2014, pp. 3606–3613. DOI: 10.1109/CVPR.2014.461.

[42] B. Courter. *Agentic Engineering: How AI Automata Will Participate in Engineering in 2025*. Blog post. 2025. URL: https://www.blakecourter.com/2025/01/20/agentic-engineering.html.

[43] R. C. Daudt. *Onera Satellite Change Detection Dataset*. Satellite change detection dataset. 2018. URL: https://rcdaudt.github.io/oscd/.

[44] T. Dettmers et al. “QLoRA: Efficient Finetuning of Quantized LLMs”. In: *arXiv preprint* (2023). Preprint. arXiv: 2305.14314 [cs.LG]. URL: https://arxiv.org/abs/2305.14314.

[45] D. S. Dewantara et al. “CNN with Multi Stage Image Data Augmentation Methods for Indonesia Rare and Protected Orchids Classification”. In: *2020 International Conference on Sustainable Information Engineering and Technology (ICOSICA)*. Bali, Indonesia: IEEE, 2020, pp. 1–5. DOI: 10.1109/ICOSICA49951.2020.9243174.

[46] A. Dhurandhar et al. “Explanations Based on the Missing: Towards Contrastive Explanations with Pertinent Negatives”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. Montreal, Canada: Curran Associates, 2018. DOI: 10.48550/arXiv.1802.07623. URL: https://arxiv.org/abs/1802.07623.

[47] F. Doshi-Velez and B. Kim. “Towards a Rigorous Science of Interpretable Machine Learning”. In: *arXiv preprint* (2017). arXiv: 1702.08608 [cs.LG]. URL: https://arxiv.org/abs/1702.08608.

[48] L. Du et al. “A Comprehensive Drought Monitoring Method Integrating MODIS and TRMM Data”. In: *International Journal of Applied Earth Observation and Geoinformation* 23 (2013), pp. 245–253. DOI: 10.1016/j.jag.2013.01.002.

[49] P. Duan, X. Kang, and S. Li. “Convolutional Neural Network for Natural Color Visualization of Hyperspectral Images”. In: *IGARSS 2019 - 2019 IEEE International Geoscience and Remote Sensing Symposium*. Yokohama, Japan: IEEE, Aug. 2019, pp. 3372–3375. DOI: 10.1109/IGARSS.2019.8898536.

[50] P. Duan et al. “Multichannel Pulse-Coupled Neural Network-Based Hyperspectral Image Visualization”. In: *IEEE Transactions on Geoscience and Remote Sensing* 58.4 (2020), pp. 2444–2456. ISSN: 1558-0644. DOI: 10.1109/TGRS.2019.2953188.

[51] J. Eckhard et al. “Outdoor Scene Reflectance Measurements Using a Bragg-Grating-Based Hyperspectral Imager”. In: *Applied Optics* 54.13 (2015), pp. D15–D24. DOI: 10.1364/AO.54.000D15. URL: http://www.ugr.es/local/colorimg/UGR_hyperspectral_image_database.html.

[52] T. Eckhard et al. “Outdoor Scene Reflectance Measurements Using a Bragg-Grating-Based Hyperspectral Imager”. In: *Applied Optics* 54.13 (2015), pp. D15–D24. DOI: 10.1364/AO.54.000D15.

[53] European Space Agency (ESA). *Sentinel-2 MSI User Guide*. Online Documentation. Frascati, Italy, 2023. URL: https://sentinels.copernicus.eu/web/sentinel/user-guides/sentinel-2-msi/resolutions/spectral.

[54] European Space Agency (ESA). *ESA Sentinel Application Platform v2.0.2*. Software Documentation. Frascati, Italy. URL: http://step.esa.int.

[55] FaostatDocs. *Food and Agriculture Data*. (accessed 7 March 2022). 2022. URL: http://www.fao.org/faostat/.

[56] S. Feng et al. “Research on Elevator Intelligent Monitoring and Grading Warning System”. In: *2021 IEEE International Conference on Computer Science, Electronic Information Engineering and Intelligent Control Technology (CEI)*. Fuzhou, China: IEEE, Sept. 2021, pp. 145–148. DOI: 10.1109/CEI52496.2021.9574579.

[57] A. Fernández et al. “SMOTE for Learning from Imbalanced Data: Progress and Challenges, Marking the 15-year Anniversary”. In: *Journal of Artificial Intelligence Research* 61 (2018), pp. 863–905. DOI: 10.1613/jair.1.11192.

[58] B. Fornberg. “Improving the Accuracy of the Trapezoidal Rule”. In: *SIAM Review* 63.1 (2021), pp. 167–180. DOI: 10.1137/18M1229353.

[59] S. Gao et al. “Superpixelwise PCA Based Data Augmentation for Hyperspectral Image Classification”. In: *Multimedia Tools and Applications* (2024), pp. 1–21. DOI: 10.1007/s11042-024-18911-8.

[60] M. Garana, M. A. Veganzones, and B. Ayerdi. *Hyperspectral Remote Sensing Scenes*. Online Dataset. Bilbao, Spain. URL: https://www.ehu.eus/ccwintco/index.php/Hyperspectral_Remote_Sensing_Scenes.

[61] D. Garreau and U. von Luxburg. “Explaining the Explainer: A First Theoretical Analysis of LIME”. In: *arXiv preprint* (2020). arXiv: 2001.03447 [cs.LG]. URL: https://arxiv.org/abs/2001.03447.

[62] F. Gascon et al. “Copernicus Sentinel-2A Calibration and Products Validation Status”. In: *Remote Sensing* 9.6 (2017). ISSN: 2072-4292. DOI: 10.3390/rs9060584. URL: https://www.mdpi.com/2072-4292/9/6/584.

[63] L. A. Gatys, A. S. Ecker, and M. Bethge. “Texture Synthesis Using Convolutional Neural Networks”. In: *Advances in Neural Information Processing Systems*. Vol. 28. 2015, pp. 262–270. DOI: 10.5555/2969239. URL: https://proceedings.neurips.cc/paper/2015/file/a5e00132373a7031000fd987a3c9f87b-Paper.pdf.

[64] A. Gavade and V. Rajpurohit. “Multi Spectral Super Resolution and Image Quality Assessment Comparative Analysis”. In: IEEE, 2016.

[65] N. Gilboy et al. *Emergency Severity Index (ESI) - A Triage Tool for Emergency Department Care, Implementation Handbook, Version 4*. Available online. Emergency Nurses Association (ENA). Schaumburg, IL, USA, 2020. URL: https://www.ena.org/docs/default-source/education-document-library/triage/esi-implementation-handbook-2020.pdf.

[66] M. Graña, M. Veganzons, and B. Ayerdi. *Hyperspectral Remote Sensing Scenes*. Online hyperspectral image dataset. 2021. URL: https://www.ehu.eus/ccwintco/index.php/Hyperspectral_Remote_Sensing_Scenes.

[67] M. Grootendorst. “BERTopic: Neural Topic Modeling with a Class-Based TF-IDF Procedure”. In: *arXiv preprint* (2022). arXiv: 2203.05794 [cs.CL]. URL: https://arxiv.org/abs/2203.05794.

[68] N. Hama, M. Mase, and A. B. Owen. “Deletion and Insertion Tests in Regression Models”. In: *arXiv preprint* (2022). arXiv: 2205.12423 [stat.ME]. URL: https://arxiv.org/abs/2205.12423.

[69] M. Hamadache, J. Jung, J. Park, et al. “A Comprehensive Review of Artificial Intelligence-Based Approaches for Rolling Element Bearing PHM: Shallow and Deep Learning”. In: *JMST Advances* 1 (2019), pp. 125–151. DOI: 10.1007/s42791-019-0016-y.

[70] B. Hanganu, L. A. Radu, and A. Băicoianu. “Machine Learning for Condition Monitoring: Latest Trend and Review”. In: *Proceedings of the 35th International Business Information Management Association Conference (IBIMA)*. WOS:000661489803090. IBIMA Publishing, 2020. ISBN: 978-0-9998551-4-0. URL: https://ibima.org/accepted-paper/machine-learning-for-condition-monitoring-latest-trend-and-review/.

[71] E. Heitz et al. “A Sliced Wasserstein Loss for Neural Texture Synthesis”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, pp. 9442–9451. DOI: 10.1109/CVPR46437.2021.00929.

[72] H. Henao et al. “Trends in Fault Diagnosis for Electrical Machines: A Review of Diagnostic Techniques”. In: *IEEE Industrial Electronics Magazine* 8.2 (2014), pp. 31–42. DOI: 10.1109/MIE.2013.2287651.

[73] A. Hernandez, A. Vilalta, and F. Moreno-Noguer. “Neural Cellular Automata Manifold”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021, pp. 10020–10028.

[74] W. S. Hong, A. D. Haimovich, and A. R. Taylor. *Predicting Hospital Admission at Emergency Department Triage Using Machine Learning*. Dataset available starting with 6 Sep 2018. 2018. URL: https://github.com/yaleemmlc/admissionprediction.

[75] W. S. Hong, A. D. Haimovich, and R. A. Taylor. “Predicting Hospital Admission at Emergency Department Triage Using Machine Learning”. In: *PLOS ONE* 13.7 (2018), e0201016. DOI: 10.1371/journal.pone.0201016.

[76] K. Horibe, K. Walker, and S. Risi. “Regenerating Soft Robots through Neural Cellular Automata”. In: *arXiv preprint* abs/2102.02579 (2021). arXiv: 2102.02579 [cs.RO]. URL: https://arxiv.org/abs/2102.02579.

[77] A. Houdard et al. “A Generative Model for Texture Synthesis based on Optimal Transport Between Feature Distributions”. In: *Journal of Mathematical Imaging and Vision* 65.1 (2022), pp. 4–28. DOI: 10.1007/s10851-022-01108-9.

[78] O. Ivanov et al. “Improving ED Emergency Severity Index Acuity Assignment Using Machine Learning and Clinical Natural Language Processing”. In: *Journal of Emergency Nursing* 47.2 (2021), 265–278.e7. DOI: 10.1016/j.jen.2020.11.001.

[79] M. Ivanovici. *AI4AGRI Summer School 2025 - EO Big Data for Agriculture*. July 2025. DOI: 10.5281/zenodo.16456370. URL: https://doi.org/10.5281/zenodo.16456370.

[80] M. Ivanovici and N. Richard. "Fractal Dimension of Color Fractal Images". In: *IEEE Transactions on Image Processing* 20.1 (2011), pp. 227–235. DOI: 10.1109/TIP.2010.2059032.

[81] Kaggle. *High-Quality Public Datasets*. (accessed 7 March 2022). 2022. URL: https://www.kaggle.com/.

[82] X. Kang et al. "Hyperspectral Anomaly Detection with Attribute and Edge-Preserving Filters". In: *IEEE Transactions on Geoscience and Remote Sensing* 55.10 (2017), pp. 5600–5611. DOI: 10.1109/TGRS.2017.2710146.

[83] X. Kang et al. "Decolorization-Based Hyperspectral Image Visualization". In: *IEEE Transactions on Geoscience and Remote Sensing* 56.8 (2018), pp. 4346–4360. DOI: 10.1109/TGRS.2018.2815588.

[84] V. Kannan, T. Zhang, and H. Li. "A Review of the Intelligent Condition Monitoring of Rolling Element Bearings". In: *Machines* 12.7 (2024). DOI: 10.3390/machines12070484.

[85] J. Kari. "Theory of Cellular Automata: A Survey". In: *Theoretical Computer Science* 334.1 (2005), pp. 3–33. DOI: 10.1016/j.tcs.2004.11.021. URL: https://www.sciencedirect.com/science/article/pii/S030439750500054X.

[86] A. Kerestely, A. Băicoianu, and R. Bocu. "A Research Study on Running Machine Learning Algorithms on Big Data with Spark". In: *Knowledge Science, Engineering and Management*. Cham, Switzerland: Springer, 2021, pp. 307–318. ISBN: 978-3-030-82136-4. DOI: 10.1007/978-3-030-82136-4_25.

[87] H. A. Khan et al. "Saliency Based Visualization of Hyper-Spectral Images". In: *2015 IEEE International Geoscience and Remote Sensing Symposium (IGARSS)*. Milan, Italy: IEEE, July 2015, pp. 1096–1099. DOI: 10.1109/IGARSS.2015.7325979.

[88] M. Khanjary. "Cellular Learning Automata: Review and Future Trend". In: *Computational Vision and Bio-Inspired Computing: Proceedings of ICCVBIC 2021*. Springer, 2022, pp. 229–238.

[89] U. Kirsch. *Structural Optimization: Fundamentals and Applications*. Berlin, Heidelberg: Springer, 1993. ISBN: 978-3-540-55919-1. DOI: 10.1007/978-3-642-84845-2.

[90] K. Korus, M. Salamak, and M. Jasiński. "Optimization of Geometric Parameters of Arch Bridges Using Visual Programming FEM Components and Genetic Algorithm". In: *Engineering Structures* 241 (2021), p. 112465. DOI: 10.1016/j.engstruct.2021.112465.

[91] B. Krawczyk. “Learning from Imbalanced Data: Open Challenges and Future Directions”. In: *Progress in Artificial Intelligence* 5.4 (2016), pp. 221–232. DOI: 10.1007/s13748-016-0094-0.

[92] A. Krizhevsky. “One Weird Trick for Parallelizing Convolutional Neural Networks”. In: *arXiv preprint* (2014). DOI: 10.48550/arXiv.1404.5997. arXiv: 1404.5997 [cs.NE]. URL: https://arxiv.org/abs/1404.5997.

[93] A. Lambora, K. Gupta, and K. Chopra. “Genetic Algorithm - A Literature Review”. In: *2019 International Conference on Machine Learning, Big Data, Cloud and Parallel Computing (COMITCon)*. IEEE, 2019, pp. 380–384. DOI: 10.1109/COMITCon.2019.8862255.

[94] R. Lasaponara et al. “On the Use of Sentinel-2 NDVI Time Series and Google Earth Engine to Detect Land-Use/Land-Cover Changes in Fire-Affected Areas”. In: *Remote Sensing* 14.19 (2022), p. 4723. DOI: 10.3390/rs14194723.

[95] S. Le Moan et al. “A Constrained Band Selection Method Based on Information Measures for Spectral Image Color Visualization”. In: *IEEE Transactions on Geoscience and Remote Sensing* 49.12 (2011), pp. 5104–5115. DOI: 10.1109/TGRS.2011.2153864.

[96] P. T.-W. Lee, O. K. Kwon, and X. Ruan. “Sustainability Challenges in Maritime Transport and Logistics Industry and Its Way Ahead”. In: *Sustainability* 11.5 (2019). DOI: 10.3390/su11051331.

[97] C. Lessmeier et al. “Condition Monitoring of Bearing Damage in Electromechanical Drive Systems by Using Motor Current Signals of Electric Motors: A Benchmark Data Set for Data-Driven Classification”. In: University of Paderborn, 2016. URL: https://www.kaggle.com/datasets/brjapon/cwru-bearing-datasets.

[98] D. Liao, S. Chen, and Y. Qian. “Visualization of Hyperspectral Images Using Moving Least Squares”. In: *2018 24th International Conference on Pattern Recognition (ICPR)*. Beijing, China: IEEE, Aug. 2018, pp. 2851–2856. DOI: 10.1109/ICPR.2018.8545794.

[99] D. N. Lipe et al. “A Modified Emergency Severity Index Level is Associated with Outcomes in Cancer Patients with COVID-19”. In: *The American Journal of Emergency Medicine* 54 (2022), pp. 111–116. DOI: 10.1016/j.ajem.2022.02.002.

[100] A. Lowe, N. Harrison, and A. P. French. “Hyperspectral Image Analysis Techniques for the Detection and Classification of the Early Onset of Plant Disease and Stress”. In: *Plant Methods* 13.1 (2017), p. 80. DOI: 10.1186/s13007-017-0233-z.

[101] R. I. Luca, A. Băicoianu, and I. C. Plajer. “Spectral Image Data Aggregation for Multisource Data Augmentation”. In: *European Journal of Remote Sensing* 58.1 (2025), p. 2492295. DOI: 10.1080/22797254.2025.2492295. URL: https://doi.org/10.1080/22797254.2025.2492295.

[102] S. Lundberg and S.-I. Lee. “A Unified Approach to Interpreting Model Predictions”. In: *arXiv preprint* (2017). arXiv: 1705.07874 [cs.AI]. URL: https://arxiv.org/abs/1705.07874.

[103] S. M. Lundberg and S.-I. Lee. “A Unified Approach to Interpreting Model Predictions”. In: *Advances in Neural Information Processing Systems 31 (NeurIPS)*. Long Beach, CA, USA: Curran Associates, Inc., 2017, pp. 4765–4774. DOI: 10.5555/3295222.3295230. URL: https://proceedings.neurips.cc/paper/2017/file/8a20a8621978632d76c43dfd28b67767-Paper.pdf.

[104] D. A. Lupșa-Tătaru et al. *SEEN Pilot course on entrepreneurship*. Aug. 2025. DOI: 10.5281/zenodo.16983067. URL: https://doi.org/10.5281/zenodo.16983067.

[105] M. Magnusson et al. “Creating RGB Images from Hyperspectral Images Using a Color Matching Function”. In: *IGARSS 2020 - 2020 IEEE International Geoscience and Remote Sensing Symposium*. Waikoloa, HI, USA: IEEE, Sept. 2020, pp. 2045–2048. DOI: 10.1109/IGARSS39084.2020.9324567.

[106] L. Majercsik, A. Băicoianu, and I. C. Plajer. “Crop-Specific Effects of Data Augmentation on Classification Accuracy in Sentinel-2 Imagery”. In: *2025 18th International Conference on Engineering of Modern Electric Systems (EMES)*. IEEE, 2025, pp. 1–6. DOI: 10.1109/EMES65692.2025.11045604.

[107] P. Manousopoulos, V. Drakopoulos, and T. Theoharis. “Curve Fitting by Fractal Interpolation”. In: *Transactions on Computational Science I*. Ed. by M. L. Gavrilova and C. J. K. Tan. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 85–103. DOI: 10.1007/978-3-540-79299-4_4.

[108] P. Manousopoulos, V. Drakopoulos, and T. Theoharis. “Curve fitting by fractal interpolation”. In: *Transactions on Computational Science I*. Springer, 2008, pp. 85–103.

[109] M. B. Marinov et al. “Smart Multisensor Node for Remote Elevator Condition Monitoring”. In: *2020 21st International Symposium on Electrical Apparatus & Technologies (SIELA)*. Bourgas, Bulgaria: IEEE, 2020, pp. 1–4. DOI: 10.1109/SIELA49118.2020.9167049.

[110] R. J. Martin et al. “XAI-Powered Smart Agriculture Framework for Enhancing Food Productivity and Sustainability”. In: *IEEE Access* 12 (2024), pp. 168412–168427. DOI: 10.1109/ACCESS.2024.3492973.

[111] M. May. “Fractal Image Compression”. In: *American Scientist* 84.5 (1996).

[112] M. Mehmood et al. “Remote Sensing Image Classification: A Comprehensive Review and Applications”. In: *Mathematical Problems in Engineering* 2022 (2022), pp. 1–24. ISSN: 1024-123X. DOI: 10.1155/2022/5880959.

[113] U. Meier et al. *Growth Stages of Mono- and Dicotyledonous Plants: BBCH Monograph*. Tech. rep. Digitized edition (E-Paper) available from JKI. Berlin / Boston: Blackwell Wissenschafts-Verlag: Julius Kühn-Institut (Federal Biological Research Centre for Agriculture and Forestry), 1997. URL: https://www.julius-kuehn.de/en/bbch/.

[114] MIT Vision and Modeling Group. *VisTex Database*. 1995. URL: https://vismod.media.mit.edu/pub/VisTex/.

[115] A. Mordvintsev and E. Niklasson. "μNCA: Texture Generation with Ultra-Compact Neural Cellular Automata". In: *arXiv preprint* (2021). arXiv: 2111.13545 [cs.LG]. URL: https://arxiv.org/abs/2111.13545.

[116] A. Mordvintsev and E. Niklasson. "μNCA: Texture Generation with Ultra-Compact Neural Cellular Automata". In: *arXiv preprint* abs/2111.13545 (2021). arXiv: 2111.13545 [cs.CV]. URL: https://arxiv.org/abs/2111.13545.

[117] A. Mordvintsev et al. "Differentiable Image Parameterizations". In: *Distill* (2018). DOI: 10.23915/distill.00012. URL: https://distill.pub/2018/differentiable-parameterizations.

[118] A. Mordvintsev et al. "Growing Neural Cellular Automata". In: *Distill* (2020). DOI: 10.23915/distill.00023. URL: https://distill.pub/2020/growing-ca.

[119] J. Nalepa et al. "Estimating Soil Parameters From Hyperspectral Images: A Benchmark Dataset and the Outcome of the HYPERVIEW Challenge". In: *IEEE Geoscience and Remote Sensing Magazine* 12.3 (2024), pp. 35–63. DOI: 10.1109/MGRS.2024.3394040.

[120] S. Nandi, H. A. Toliyat, and X. Li. "Condition Monitoring and Fault Diagnosis of Electrical Motors—A Review". In: *IEEE Transactions on Energy Conversion* 20.4 (2005), pp. 719–729. DOI: 10.1109/TEC.2005.847955.

[121] M. A. Navascués. "Reconstruction of Sampled Signals with Fractal Functions". In: *Acta Applicandae Mathematicae* 110 (2010).

[122] T.-C. Negura, I. Popa, and A. Băicoianu. "Race Line Optimization with Metaheuristics". In: *Proceedings of the International Symposium on Systems and Control Sciences (ISCS 2025)*. Accepted for publication, in press. 2025. URL: https://scs.etti.tuiasi.ro/isscs2025/technical_program.html#A2.

[123] Z. Nemes et al. "The Study of the Potato's Life-Cycle Phases Important to the Increase of the Individual Variability". In: *Analele Universității din Oradea, Fascicula Biologie* 15 (2008), pp. 60–63. ISSN: 1222-5119.

[124] E. Niklasson et al. "Self-Organising Textures". In: *Distill* (2021). DOI: 10.23915/distill.00027.003. URL: https://distill.pub/selforg/2021/textures/.

[125] C. Oh, S. Han, and J. Jeong. “Time-Series Data Augmentation Based on Interpolation”. In: *Proceedings of FNC/MobiSPC*. 2020. URL: https://api.semanticscholar.org/CorpusID:222112073.

[126] I. O. Olalere, M. Dewa, and B. Nleya. “Remote Condition Monitoring of Elevator’s Vibration and Acoustics Parameters for Optimised Maintenance Using IoT Technology”. In: *2018 IEEE Canadian Conference on Electrical & Computer Engineering (CCECE)*. Quebec City, Canada: IEEE, 2018, pp. 1–4. DOI: 10.1109/CCECE.2018.8447771.

[127] OpenWeather. *Visualisation of the NDVI Index on Satellite Maps: Custom Palettes for Agricultural Applications*. Online Article. 2021. URL: https://openweathermap.medium.com/visualisation-of-the-ndvi-index-on-satellite-maps-custom-palettes-for-agricultural-applications-f99b0652f991.

[128] E. Pajouheshgar et al. “DyNCA: Real-Time Dynamic Texture Synthesis Using Neural Cellular Automata”. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2023, pp. 20742–20751.

[129] J. S. Park et al. “Generative Agents: Interactive Simulacra of Human Behavior”. In: *arXiv preprint* (2023). Preprint. arXiv: 2304.03442 [cs.HC]. URL: https://arxiv.org/abs/2304.03442.

[130] I. C. Plajer, A. Băicoianu, and M. Ivanovici. “Multispectral Image Visualization and NDVI-Based Anomaly Detection”. In: *Bulletin of the Transilvania University of Brasov. Series III: Mathematics and Computer Science* (2025). Accepted for publication, in press.

[131] I. C. Plajer, A. Băicoianu, and L. Majercsik. “AI-Based Visualization of Remotely-Sensed Spectral Images”. In: *2023 International Symposium on Signals, Circuits and Systems (ISSCS)*. Conference held in Iasi, Romania. New York, NY, USA: IEEE, July 2023, pp. 1–4. DOI: 10.1109/ISSCS58449.2023.10190908. URL: https://doi.org/10.1109/ISSCS58449.2023.10190908.

[132] I. C. Plajer et al. “NDVI Computation from Hyperspectral Images”. In: *2023 13th Workshop on Hyperspectral Imaging and Signal Processing: Evolution in Remote Sensing (WHISPERS)*. IEEE, 2023, pp. 1–5. DOI: 10.1109/WHISPERS61460.2023.10431225.

[133] I. C. Plajer et al. “Multisource Remote Sensing Data Visualization Using Machine Learning”. In: *IEEE Transactions on Geoscience and Remote Sensing* 62 (2024), pp. 1–12. DOI: 10.1109/TGRS.2024.3372639.

[134] I. C. Plajer et al. “Multi-Year Multi-Crop Correlation Analysis in Brasov Area”. In: *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences* XLVIII-M-7-2025 (2025), pp. 195–200. DOI: 10.5194/isprs-

archives - XLVIII - M - 7 - 2025 - 195 - 2025. URL: https://isprs-archives.copernicus.org/articles/XLVIII-M-7-2025/195/2025/.

[135] R. P. Porfírio. "Studying the Impact of Explainable AI in Digital Agriculture Solutions". In: *Companion Publication of the 2024 Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*. San Jose, Costa Rica: ACM, 2024, pp. 43–46. DOI: 10.1145/3678884.3682046.

[136] A. B. Pour, M. Parsa, and A. M. Eldosouky, eds. *Geospatial Analysis Applied to Mineral Exploration: Remote Sensing, GIS, Geochemical, and Geophysical Applications to Mineral Resources*. Print ISBN 978-0-323-95608-6; eBook ISBN 978-0-323-95609-3. Philadelphia, PA, USA: Elsevier, 2023, p. 322. ISBN: 978-0-323-95608-6.

[137] A. Radford et al. "Language Models are Unsupervised Multitask Learners". In: (2019).

[138] I. Rădulescu, A. Băicoianu, and A. Mihai. *Conformal Transformation of Kernels: A Geometric Perspective on Text Classification*. Accepted at Facta Universitatis. 2024. arXiv: 2406.00499 [cs.LG]. URL: https://arxiv.org/abs/2406.00499.

[139] I. Rădulescu, A. Băicoianu, and A. Mihai. *Conformal Transformation of Kernels: A Geometric Perspective on Text Classification*. 2024. arXiv: 2406.00499 [cs.LG]. URL: https://arxiv.org/abs/2406.00499.

[140] A. Rai and S. Upadhyay. "A Review on Signal Processing Techniques Utilized in the Fault Diagnosis of Rolling Element Bearings". In: *Tribology International* 96 (2016), pp. 289–306. DOI: 10.1016/j.triboint.2015.12.037.

[141] M. R. Raia et al. "Condition Monitoring of Industrial Elevators Based on Machine Learning Models". In: *2023 IEEE 28th International Conference on Emerging Technologies and Factory Automation (ETFA)*. Sinaia, Romania: IEEE, 2023, pp. 1–5. DOI: 10.1109/ETFA54631.2023.10275563.

[142] S. Raubitzek and T. Neubauer. "A Fractal Interpolation Approach to Improve Neural Network Predictions for Difficult Time Series Data". In: *Expert Systems with Applications* 169 (2021), p. 114474. DOI: 10.1016/j.eswa.2020.114474.

[143] J. Rayo, R. de la Rosa, and M. Garrido. "A Hybrid Approach to Information Retrieval and Answer Generation for Regulatory Texts". In: *arXiv preprint* (2025). arXiv: 2502.16767 [cs.CL]. URL: https://arxiv.org/abs/2502.16767.

[144] C. Rodriguez-Pardo and E. Garces. "SeamlessGAN: Self-Supervised Synthesis of Tileable Texture Maps". In: *IEEE Transactions on Visualization and Computer Graphics* 29.6 (2022), pp. 2914–2925. DOI: 10.1109/TVCG.2022.3157048.

[145] P. Rodríguez-Veiga et al. "Forest Biomass Retrieval Approaches from Earth Observation in Different Biomes". In: *International Journal of Applied Earth Observation and Geoinformation* 77 (2019), pp. 53–68. DOI: 10.1016/j.jag.2018.12.008.

[146] J. E. Ruiz-Sarrio, J. A. Antonino-Daviu, and C. Martis. “Localized Bearing Fault Analysis for Different Induction Machine Start-Up Modes via Vibration Time–Frequency Envelope Spectrum”. In: *Sensors* 24.21 (2024). DOI: 10.3390/s24216935.

[147] S. Sacareanu et al. “Test-Based Virtual Methods to Explore NVH Issues of Hybrid and Electrical Vehicles”. In: *CONAT 2024 International Congress of Automotive and Transport Engineering*. Ed. by A. Chiru and D. Covaciu. Cham: Springer Nature Switzerland, 2025, pp. 232–242. ISBN: 978-3-031-77627-4.

[148] E. Salari and Z. Azimifar. “Texture Image Synthesis Using Spatial GAN Based on Vision Transformers”. In: *arXiv preprint* (2025). arXiv: 2502.01842 [cs.CV].

[149] R. Sánchez-Salmerón et al. “Machine Learning Methods Applied to Triage in Emergency Services: A Systematic Review”. In: *International Emergency Nursing* 60 (2022), p. 101109. DOI: 10.1016/j.ienj.2021.101109.

[150] R. R. Selvaraju et al. “Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization”. In: *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*. Venice, Italy: IEEE, Sept. 2017, pp. 618–626. DOI: 10.1109/ICCV.2017.74.

[151] SentinelHub. *Normalized Difference Vegetation Index (NDVI)*. Online Script Documentation. 2022. URL: https://custom-scripts.sentinel-hub.com/custom-scripts/sentinel-2/ndvi/.

[152] U. Shafi et al. “A Multi-Modal Approach for Crop Health Mapping Using Low Altitude Remote Sensing, Internet of Things (IoT) and Machine Learning”. In: *IEEE Access* 8 (2020), pp. 112708–112724. DOI: 10.1109/ACCESS.2020.3002940.

[153] H. Shao et al. “Intelligent Fault Diagnosis of Rotor-Bearing System Under Varying Working Conditions With Modified Transfer Convolutional Neural Network and Thermal Images”. In: *IEEE Transactions on Industrial Informatics* 17.5 (2021), pp. 3488–3496. DOI: 10.1109/TII.2020.3005965.

[154] P. Sharma et al. “Recent Advances in Machine Learning Research for Nanofluid-Based Heat Transfer in Renewable Energy System”. In: *Energy & Fuels* 36.13 (2022), pp. 6626–6658. DOI: 10.1021/acs.energyfuels.2c01006.

[155] M. Shehab et al. “Machine Learning in Medical Applications: A Review of State-of-the-Art Methods”. In: *Computers in Biology and Medicine* 145 (2022), p. 105458. DOI: 10.1016/j.compbiomed.2022.105458.

[156] W. Shi et al. “Direct Multi-Turn Preference Optimization for Language Agents”. In: *arXiv preprint* (2025). Preprint. arXiv: 2406.14868 [cs.CL]. URL: https://arxiv.org/abs/2406.14868.

[157] S. Y. Shin et al. “Towards Generating Executable Metamorphic Relations Using Large Language Models”. In: *Quality of Information and Communications Technology*. Cham, Switzerland: Springer, 2024, pp. 126–141. ISBN: 978-3-031-70245-7. DOI: 10.1007/978-3-031-70245-7_9.

[158] A. Singh et al. “SciRepEval: A Multi-Format Benchmark for Scientific Document Representations”. In: *Proceedings of the Conference on Empirical Methods in Natural Language Processing*. 2022. URL: https://api.semanticscholar.org/CorpusID:254018137.

[159] P. Singh et al. “Hyperspectral Remote Sensing in Precision Agriculture: Present Status, Challenges, and Future Trends”. In: *Hyperspectral Remote Sensing*. Elsevier, 2020, pp. 121–146. DOI: 10.1016/B978-0-12-819453-9.00005-3.

[160] R. P. Sishodia, R. L. Ray, and S. K. Singh. “Applications of Remote Sensing in Precision Agriculture: A Review”. In: *Remote Sensing* 12.19 (2020), p. 3136. DOI: 10.3390/rs12193136.

[161] *SNAP - Resampling Methods*. Online documentation for Sentinel Application Platform (SNAP). URL: https://step.esa.int/main/wp-content/help/versions/11.0.0/snap/org.esa.snap.snap.help/general/overview/ResamplingMethods.html.

[162] D. I. Stoica, I. Popa, and A. Băicoianu. “Automated Model Interpretation, Pattern Analysis and Insights in Explainable AI”. In: *Innovative Perspectives on Computational Intelligence and Data Science* (2025). Accepted for publication, in press.

[163] J. Stovold. “Neural Cellular Automata Can Respond to Signals”. In: *Proceedings of the Artificial Life Conference*. 2023, p. 5. DOI: 10.1162/isal_a_00567.

[164] H. Su, Q. Du, and P. Du. “Hyperspectral Image Visualization Using Band Selection”. In: *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 7.6 (2013), pp. 2647–2658. DOI: 10.1109/JSTARS.2013.2287696.

[165] Y. Sun et al. “Evolving Deep Convolutional Neural Networks for Image Classification”. In: *IEEE Transactions on Evolutionary Computation* 24.2 (2020), pp. 394–407. DOI: 10.1109/TEVC.2019.2916183.

[166] K. Szczepankiewicz et al. “Ground Truth-Based Comparison of Saliency Maps Algorithms”. In: *Scientific Reports* 13.1 (2023), p. 16887. DOI: 10.1038/s41598-023-42946-w.

[167] K. J. W. Tang et al. “Artificial Intelligence and Machine Learning in Emergency Medicine”. In: *Biocybernetics and Biomedical Engineering* 41.1 (2021), pp. 156–172. DOI: 10.1016/j.bbe.2020.12.002.

[168] R. Tang et al. “Supervised Learning with Convolutional Neural Networks for Hyperspectral Visualization”. In: *Remote Sensing Letters* 11.4 (2020), pp. 363–372. DOI: 10.1080/2150704X.2020.1730469.

[169] V. Tsagaris, V. Anastassopoulos, and G. A. Lampropoulos. “Fusion of Hyperspectral Data Using Segmented PCT for Color Representation and Classification”. In: *IEEE Transactions on Geoscience and Remote Sensing* 43.10 (2005), pp. 2365–2375. DOI: 10.1109/TGRS.2005.856106.

[170] Ö. Turgut, İ. Kök, and S. Özdemir. “AgroXAI: Explainable AI-Driven Crop Recommendation System for Agriculture 4.0”. In: *2024 IEEE International Conference on Big Data (BigData)*. Washington, DC, USA: IEEE, 2024, pp. 7208–7217. DOI: 10.1109/BigData62323.2024.10825771.

[171] D. Ulyanov, A. Vedaldi, and V. Lempitsky. “Improved Texture Networks: Maximizing Quality and Diversity in Feed-Forward Stylization and Texture Synthesis”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2017, pp. 4105–4113. DOI: 10.1109/CVPR.2017.437.

[172] A. Vântu, A. Vasilescu, and A. Băicoianu. “Medical Emergency Department Triage Data Processing Using a Machine-Learning Solution”. In: *Heliyon* 9.8 (2023), e18402. DOI: 10.1016/j.heliyon.2023.e18402.

[173] A. Vaswani et al. “Attention Is All You Need”. In: *arXiv preprint* (2023). Preprint. arXiv: 1706.03762 [cs.CL]. URL: https://arxiv.org/abs/1706.03762.

[174] I. Vlascu and A. Băicoianu. “CNN-Based Violence Detection: Evaluation and Experiments”. In: *2025 18th International Conference on Engineering of Modern Electric Systems (EMES)*. New York, NY, USA: IEEE, June 2025, pp. 1–6. DOI: 10.1109/EMES65692.2025.11045585.

[175] I. Vlascu, A. Băicoianu, and M. Ivanovici. “Improving CNN-Based Video Violence Detection”. In: *2025 International Symposium on Signals, Circuits and Systems (ISSCS)*. 2025, pp. 1–4. DOI: 10.1109/ISSCS66034.2025.11105596.

[176] H. Wang et al. “Score-CAM: Score-Weighted Visual Explanations for Convolutional Neural Networks”. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. Seattle, WA, USA: IEEE, June 2020, pp. 111–119. DOI: 10.1109/CVPRW50498.2020.00020.

[177] H.-Y. Wang, H. Li, and J.-Y. Shen. “A Novel Hybrid Fractal Interpolation-SVM Model for Forecasting Stock Price Indexes”. In: *Fractals* 27.4 (2019), pp. 1950055–18. DOI: 10.1142/S0218348X19500555.

[178] Z. Wang et al. “A Globally Interpretable Convolutional Neural Network Combining Bearing Semantics for Bearing Fault Diagnosis”. In: *IEEE Transactions on Instrumentation and Measurement* 74 (2025), pp. 1–13. DOI: 10.1109/TIM.2025.3538068.

[179] Z. Wang et al. “Image Quality Assessment: From Error Visibility to Structural Similarity”. In: *IEEE Transactions on Image Processing* 13.4 (2004), pp. 600–612. DOI: 10.1109/TIP.2003.819861.

[180] N. F. Waziralilah et al. “A Review on Convolutional Neural Network in Bearing Fault Diagnosis”. In: *MATEC Web of Conferences* 255 (2019), p. 06002. DOI: 10.1051/matecconf/201925506002.

[181] M. Weiss, F. Jacob, and G. Duveiller. “Remote Sensing for Agricultural Applications: A Meta-Review”. In: *Remote Sensing of Environment* 236 (2020), p. 111402. DOI: 10.1016/j.rse.2019.111402.

[182] E. W. Weisstein. *Elementary Cellular Automaton*. Wolfram MathWorld. 2002. URL: https://mathworld.wolfram.com/ElementaryCellularAutomaton.html.

[183] Q. Wen et al. “Time Series Data Augmentation for Deep Learning: A Survey”. In: *Proceedings of the 30th International Joint Conference on Artificial Intelligence (IJCAI)*. 2021.

[184] S. Winkler. “Vision Models and Quality Metrics for Image Processing Applications”. PhD Thesis. PhD thesis. Lausanne, Switzerland: École Polytechnique Fédérale de Lausanne (EPFL), 2001. URL: https://infoscience.epfl.ch/record/100027.

[185] S. Wolfram. *A New Kind of Science*. Wolfram Media, 2002. ISBN: 1579550088. URL: https://www.wolframscience.com.

[186] Z. Xi et al. “The Rise and Potential of Large Language Model Based Agents: A Survey”. In: *arXiv preprint* (2023). Preprint. arXiv: 2309.07864 [cs.AI]. URL: https://arxiv.org/abs/2309.07864.

[187] S. Yadav and A. Sohal. “Comparative Study of Different Selection Techniques in Genetic Algorithms”. In: *International Journal of Engineering Sciences and Mathematics* 6.3 (2017), pp. 174–180. ISSN: 2320-0294.

[188] F. Yakuwa et al. “Novel Time Series Analysis and Prediction of Stock Trading Using Fractal Theory and Time Delayed Neural Network”. In: *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics*. Vol. 1. 2003, pp. 134–141. DOI: 10.1109/ICSMC.2003.1243804.

[189] S. Yao et al. “ReAct: Synergizing Reasoning and Acting in Language Models”. In: *arXiv preprint* (2023). Preprint. arXiv: 2210.03629 [cs.CL]. URL: https://arxiv.org/abs/2210.03629.

[190] M. H. Yarmohammadian et al. “Overcrowding in Emergency Departments: A Review of Strategies to Decrease Future Challenges”. In: *Journal of Research in Medical Sciences* 22 (2017), p. 23. DOI: 10.4103/1735-1995.200277.

[191] F. Yasuma et al. *CAVE Dataset: Multispectral Image Database*. Multispectral image dataset. 2008. URL: https://www.cs.columbia.edu/CAVE/databases/multispectral/.

[192] F. Yasuma et al. “Generalized Assorted Pixel Camera: Postcapture Control of Resolution, Dynamic Range, and Spectrum”. In: *IEEE Transactions on Image Processing* 19.9 (2010), pp. 2241–2253. ISSN: 1057-7149. DOI: 10.1109/TIP.2010.2046811.

[193] D. A. Zaitsev. “k-Neighborhood for Cellular Automata”. In: *arXiv preprint* abs/ 1605.08870 (2016). arXiv: 1605.08870 [cs.DM]. URL: https://arxiv.org/abs/1605.08870.

[194] M. D. Zeiler and R. Fergus. “Visualizing and Understanding Convolutional Networks”. In: *Computer Vision – ECCV 2014*. Zurich, Switzerland: Springer, 2014, pp. 818–833. DOI: 10.1007/978-3-319-10590-1_53.

[195] B. Zhou et al. “Learning Deep Features for Discriminative Localization”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. Las Vegas, NV, USA: IEEE, June 2016, pp. 2921–2929. DOI: 10.1109/CVPR.2016.319.

[196] Y. Zhou. *Agentic AI Engineering: The Blueprint for Production-Grade AI Agents*. Medium article. July 2025. URL: https://medium.com/generative-ai-revolution-ai-native-transformation/agentic-ai-engineering-the-blueprint-for-production-grade-ai-agents-20358468b0b1.

[197] Y. Zhu, P. K. Varshney, and H. Chen. “Evaluation of ICA Based Fusion of Hyperspectral Images for Color Display”. In: *2007 10th International Conference on Information Fusion*. Quebec, Canada: IEEE, July 2007, pp. 1–7. DOI: 10.1109/ICIF.2007.4408004.