


Environment Perception, Representation and Understanding from Aerial Perspective – Sergiu Nedeveschi and Florin Oniga


The work addresses several key challenges in **AI-based aerial scene understanding and reconstruction**, combining new datasets, complexity metrics, geometric reconstruction methods, depth estimation, and semantic segmentation.

A major result is **ClaraVid**, a synthetic multimodal aerial dataset designed for joint semantic and geometric reconstruction from UAV imagery. It contains nearly 17,000 high-resolution frames with RGB images, metric depth, panoptic segmentation, dynamic-object masks, point clouds, and camera calibration data. The dataset also introduces the **Delentropic Scene Profile (DSP)** to quantify scene complexity and relate it to reconstruction accuracy; the work was published at **ICCV 2025 (A*)** [1].



Overview of ClaraVid Dataset. (a) Multi-viewpoint UAV acquisition, (b) High-fidelity, diverse environments and (c) Multimodal ground truth.

Building on this idea, a new **complexity-aware Level-of-Detail method for 3D Gaussian Splatting** was developed, allocating Gaussian primitives according to local image complexity. The approach improves reconstruction efficiency and quality across multiple LoD levels and was submitted to **NeurIPS 2026**.




Complexity vs. Reconstruction Quality Correlation. Complexity and quality correlate both in the 2D image plane and in the 3D object space. The complexity color is inverted for visualization.

For **monocular metric depth estimation from UAV imagery**, the TerraDepth approach incorporates coarse terrain information from TanDEM-X directly into the depth-estimation pipeline. By conditioning the metric depth estimation on terrain geometry, it improves absolute-scale reconstruction while preserving image-based local detail [2].



Significant advances were also achieved in **aerial semantic segmentation**. The UAVid dataset was thoroughly re-annotated and extended into **UAVid++**, improving label quality, reducing class imbalance, and making the benchmark more discriminative [3]. In parallel, a DINO-based segmentation pipeline with feature pyramids, multi-scale training, and overlapping-window inference substantially improved both small-object segmentation and cross-domain generalization, reaching up to **85.11% mIoU on UAVid++** [4].



The team has concluded a cooperation agreement with **Start Business SRL** for the development of a small cargo transport drone. The team is responsible for developing the software component.

The contributions strengthen UAV perception through better datasets, scene-complexity modeling, geometry-aware depth estimation, and highly generalizable foundation-model-based semantic segmentation.

[1] R. Beche, S. Nedeveschi, **ClaraVid**: A Holistic Scene Reconstruction Benchmark From Aerial Perspective With Delentropy-Based Complexity Profiling, *ICCV*, October 2025 (A*)

[2] Florea, H. and Nedeveschi, S., 2025. **Tandepth**: Leveraging global dems for metric monocular depth estimation in uavs. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 18, pp.5445-5459 (Q1).

[3] V. Chiciudean, S. Nedeveschi and F. Oniga, "UAVid++: Higher-Quality Labels and Expanded Semantic Taxonomy for Aerial Semantic Segmentation," in *IEEE Transactions on Geoscience and Remote Sensing (TGRS)*, vol. 64, pp. 5634816-5634816, 2026 (Q1).

[4] V. Chiciudean, S. Nedeveschi and F. Oniga, "Boosting Small Object Segmentation and Generalization of DINO-Based Vision Transformers in Aerial Imagery" - currently presented at *IEEE Intelligent Transportation Systems Conference (ITSC)* 2026.

Intelligent Space Surveillance – Radu Dănescu

Within the **HRIA project**, the TUCN team, together with the Astronomical Observatory of Cluj-Napoca, contributed to ESA's **Cluster II satellite reentry observation campaigns**, participating in the monitoring of the **Rumba, Samba and Tango** satellites. These campaigns, coordinated by the ESA Space Debris Office, aimed to provide accurate optical observations for assessing and supporting the controlled atmospheric reentries.

The team used high-accuracy optical instrumentation capable of tracking objects from LEO to HEO and developed dedicated **HRIA software for automatic processing of image sequences**, including target detection, astrometric calibration and generation of tracklets in equatorial coordinates. The main technical challenges were the very low signal-to-noise ratio at distances above 100,000 km, star streaks caused by object tracking, and the rolling-shutter characteristics of CMOS cameras.

The developed solution uses **image stacking to enhance faint targets, automatic detection of star streaks for calibration, object detection in stacked images, and conversion of image coordinates into RA/Dec coordinates**. The methodology was published in *Sensors* in 2026.

The campaigns produced particularly strong results: for **Rumba (2025)**, 55 tracklets were generated during 13 observation nights, while for **Samba and Tango (2026)**, 50 tracklets were generated during 12 nights. Following hardware and software improvements, the 2026 observations achieved a **100% ESA acceptance rate**, with the satellites successfully tracked up to approximately **one hour before reentry**. ESA explicitly acknowledged the high value of the contributions and expressed interest in continued collaboration and scientific follow-up.

The work demonstrates a successful HRIA outcome combining specialized optical hardware and original automated image-processing software, validated in demanding real-world ESA space-surveillance campaigns.

[1] R. Danescu, V. Turcu, "[Automatic Data Reduction of Image Sequences Acquired in Object Tracking Mode for Detection and Position Measurement of Faint Orbital Objects](#)", *Sensors*, Vol. 26, No. 5, 2026, Art. No. 1628.


Intelligent Applications for Smart Cities and the Environment – Cioară Tudor

The HRIA results significantly strengthened UTCN's **research capacity, European visibility, and technology transfer in AI for energy systems**. Building on expertise developed within HRIA, UTCN secured two new **Horizon Europe Innovation Action projects, SMART-ESS and GenAISIS**. The projects focus on Physics-Informed Neural Networks, Digital Twins, Trustworthy AI, Federated Learning, Generative AI, Foundation Models, and efficient AI deployment on resource-constrained devices.

UTCN also strengthened its European position through participation in the **AI.grids initiative**, a consortium of 54 organizations developing a European foundation model for energy-grid operations, with UTCN represented on its Supervisory Board and contributing to scalable AI training, GenAI/MLOps, and HPC/GPU infrastructures.

The work also produced **high-impact Q1 publications [1-3]** in AI and energy, while supporting **technology transfer to Romanian industry** through collaborations with Rebeldot, E-INFRA, Distribuție Energie Electrică Romania, Transelectrica, and Yellow Grid. These activities aim at validating HRIA-developed AI technologies up to **TRL 6** and integrating Romanian industrial partners into future Horizon Europe projects.

The proposed architecture of a Agentic DTs showing the FMs support the energy management phases.



The team created a strong basis for new European projects, strategic AI-for-energy partnerships, scientific visibility, and transfer of advanced AI technologies toward real-world energy applications.

[1] Gabriel Antonesi, Tudor Cioara, Ionut Anghel, Vasilis Michalakopoulos, Elissaios Sarmas, Liana Todorean, A systematic review of transformers and large language models in the energy sector: towards agentic digital twins, *Applied Energy*, Vol. 401, Part A, 2025, 126670, ISSN 0306-2619, <https://doi.org/10.1016/j.apenergy.2025.126670>.

IF 11, Q1

[2] G. Antonesi, T. Cioara, V. Michalakopoulos, I. Papias, I. Anghel, E. Sarmas, Energy forecasting under missing data: Comparative evaluation of augmented representations and decoder-only time-series imputation, *Energy and AI*, Volume 25, 2026, 100780, ISSN 2666-5468, <https://doi.org/10.1016/j.egyai.2026.100780>. **IF 9.6, Q1**

[3] R. Martin, T. Cioara, I. Anghel, G. Arcas, Performance evaluation of large language models for automated knowledge graph generation, *Machine Learning with Applications*, Vol. 24, 2026, 100923, ISSN 2666-8270, <https://doi.org/10.1016/j.mlwa.2026.100923> **IF 6.1, Q1**.

AI-based energy demand forecasting and management – Rodica Potolea

In the energy domain, the work targets **smart household energy management** through demand forecasting and optimized appliance scheduling. The HeySF framework achieved a MAPE of **3.53%** for scalable demand forecasting, while ContrastRelevance addressed strongly imbalanced usage patterns. A closed-loop demand-side management system combining explainable AI and multisensor fusion reduced forecasting errors by **29%** and peak demand by up to **36.36%**, with no net energy loss [1-2].

Additional work introduced **TGSP**, a self-supervised test-time refinement method for improving object-detection confidence without retraining or extra labeled data, extending the applicability of the research toward robust AI inference under limited annotation conditions.

The contributions advance trustworthy and explainable AI through neuro-symbolic reasoning, knowledge-based NLP, robust forecasting, multisensor fusion, and self-supervised inference, with applications in fact-checking and smart energy management.

[1] E. Hadaarau, R. Portase, and R. Potolea, "Energy Foresight: Predicting Daily Demand with Combined Techniques," in 2025 IEEE 21st International Conference on Intelligent Computer Communication and Processing (ICCP), Cluj-Napoca, Romania, 2025.


[2] R. L. Portase and R. Potolea, "Robust explainable forecasting via heterogeneous sensor fusion with applications in residential peak load management," *Energy and AI*, vol. 24, p. 100764, 2026. [Q1 Journal]

AI methods for biomedical applications –Anca Marginean si Adrian Groza

The HRIA contributions presented in the document focus mainly on **AI methods for biomedical applications**, especially ophthalmology, together with more general methods for improving the reliability and interpretability of machine-learning systems.


A first direction concerns **AI-assisted biomedical diagnosis**. In collaboration with UMF “Iuliu Hațieganu”, work addressed bacterial detection in complex samples and the use of AI to compensate for matrix effects, biological variability and sensor heterogeneity. In ophthalmology, the MIRAGE model was adapted for **drusen and retinal-fluid segmentation in OCT**, enabling slice-by-slice analysis, 3D reconstruction and quantitative estimation of fluid volumes. [1]

A major result is the development of an **OCT-specific foundation model based on DINOv2**, pretrained on approximately 47,000 OCT images. The domain-adapted model significantly outperformed the generic DINOv2 baseline, reaching a macro-F1 of **0.775 versus 0.662** on an external DME severity dataset, while requiring substantially less self-supervised pretraining data than RETFound. (<https://github.com/dariaungureanu/retino-oct-dinov2>) [2]




A second research direction targets **robust and trustworthy machine learning**. A methodology based on Monte-Carlo dropout was proposed for detecting potentially corrupted labels, estimating the amount of data requiring re-annotation and retraining a “self-healed” model. The method is designed to be task-, dataset- and architecture-independent and was validated on both image and text classification problems.[3]

For surgical video understanding, the **OphSurgeon** hierarchical framework was developed for ophthalmic surgical-phase recognition. By exploiting surgery-type context, class-balanced sampling and feature-space token dropout, it improved Phase Macro-F1 by about **13.2% relative to the baseline**. [4]



Finally, the work also explores **interpretable diagnostic assistance**, combining ML predictions with the Toulmin argumentation model. Biomarker evidence, medical reasoning, confidence assessment and retrieval of similar cases are integrated using models such as MedGemma and MedSigLip to support a more critical and informed evaluation of automatically generated diagnoses.[5]




The contributions demonstrate HRIA results in medical AI centered on foundation models, robust learning, uncertainty handling, multimodal reasoning and explainable decision support, with several works already published, presented or submitted for publication.


- [1] Ana-Maria Tataru, Andreea Cernat, Mihaela Tertis, Anda Fălăuş, Adrian Groza, Anca Mărginean, Cecilia Cristea , Matrix effect resolution in artificial intelligence-driven bacteria detection, *TRAC Trends in Analytical Chemistry*, vol 202, 2026
- [2] Ioana Damian, Simona Delia Nicoara, Anca Marginean PP-RET-238 Decoding macular oedema: causes, consequences, and clues from OCT, [40th World Ophthalmology Congress in Prague](#).
- [3] Adrian Ciu, Anca Marginean, Beyond Re-Annotation: Task- and Modality-Agnostic Label Noise Detection and Self-Healing via Prediction Evolution Certainty, submitted to *IEEE Access*
- [4] Cornel Alexandru Badea, Anca Marginean, OphSurgeon: Hierarchical Probing of Surgical Foundation Models for Ophthalmic Phase Recognition, submitted to *IEEE Journal of Biomedical and Health Informatics*
- [5] Anca Marginean, Adrian Groza, [From ML Predictions to Informed Diagnostic Assistance Using the Toulmin Model of Argumentation](#) , Workshop [EUt+ Dats and Stats 2026](#),

Medical Image processing – Tiberiu Marița

Development of deep learning solutions for medical imaging, dedicated to the segmentation of anatomical structures and computer-aided diagnosis using CT, MRI, X-ray, and OCT images. Models for retinal vascular network segmentation, thoracic pathology detection, and automated scoliosis assessment through Cobb angle measurement were improved. The use of domain adaptation learning (DAL) techniques for the recognition of liver and pancreatic tumors improved the performance of classification models. AI-based solutions are being developed for liver lesion characterization, fibrosis assessment, treatment response monitoring, and the diagnosis of portal hypertension.



Development of an automated method for segmenting the ventricles and myocardium in cardiac MRI images by reconfiguring the MIST architecture with attention gates, a contour-oriented loss function, and Mamba modules in the decoder. The proposed method improved segmentation accuracy, achieving a Mean DSC of 90.44% on the ACDC test set, representing an approximately 1.5% improvement over the baseline model. The integration of Mamba modules and attention mechanisms enabled these improvements with a limited impact on model complexity and inference time.



The team cooperate with „Iuliu Hațieganu” University of Medicine And Pharmacy Cluj-Napoca, Department of Medical Imaging and Regional Institute of Gastroenterology and Hepatology “Prof.dr. O. Fodor” (Prof. dr. Monica Platon and Lecturer dr. Mihai Socaciu) in the following research teams:

- AI-assisted methods for classifying focal liver lesions based on abbreviated MRI protocols
- Automated, AI-assisted methods for classifying liver fibrosis based on CT/MRI/ultrasound images
- Automated, AI-assisted methods in the diagnosis of portal hypertension
- Automated, AI-assisted methods for the evaluation of the response to the treatment of a liver lesions
- AI model for estimating Hepatic Venous Pressure Gradient (HVPG) and prognosis for patients with Advanced Chronic Liver Disease (ACLD)
- CEUS analysis based on radiomics and intensity parameters for the prognosis of patients with advanced compensated liver diseases

[1] D.I. Stan, M. B. Jurcă, R.D. Brehar, “Segmentation of the Retinal Vascular Network and Biomarker Quantification in OCTA Imaging”, *2025 IEEE 21st ICCP*, 16-18 Oct. 2025

[2] I. Oltean, T. Marița, “Inception_RetNet: Pioneering Diabetic Retinopathy Detection with Dual Attention Innovation”, *2025 IEEE 21st ICCP*, 16-18 Oct. 2025.

[3] D.E. Mihăilă, R.D. Brehar, “Automatic Scoliosis Assessment via Cobb Angle Estimation using I-MedSAM-Based Instance Segmentation”, *2025 IEEE 21st ICCP*, 16-18 Oct. 2025.

[4] R. Andro, C.C. Vancea, “Architectural Enhancements to YOLOv12 for Improving Detection of Thoracic Pathologies in Chest X-rays”, *2025 IEEE 21st ICCP*, 16-18 Oct. 2025.

Natural Language Processing for trustworthy AI – Rodica Potolea

The HRIA contributions cover two main directions: **Natural Language Processing for trustworthy AI** and **AI-based energy demand forecasting and management**.

In NLP, the long-term objective is a **robust and explainable Romanian fact-checking system** combining structured and unstructured knowledge. Key contributions include MorphNLI for stepwise natural language inference, multilingual attention-head pruning for more efficient and interpretable models, knowledge-graph integration for evidence retrieval, analysis of social bias in generative AI in relation to the EU AI Act, and a neuro-symbolic reasoning system that combines a language model with symbolic verification to reduce belief bias [1-3].

[1] Negru, V., Vacareanu, R., Lemnaru, C., Surdeanu, M., & Potolea, R. (2025). MorphNLI: A Stepwise Approach to Natural Language Inference Using Text Morphing. In Findings of the Association for Computational Linguistics: NAACL 2025 (pp. 6953–6968). Association for Computational Linguistics.


[2] Anghelus, A.G., Negru, V.A., Lemnaru, C., & Potolea, R. (2025). Effects of Attention Head Pruning on Encoder-only Language Models for Multilingual Recipe Classification. In 2025 International Conference on Speech Technology and Human-Computer Dialogue (SpeD) (pp. 29-33).

[3] Butas, R., Lapusan, A., Lemnaru, C., & Potolea, R. (2026). TUCNLP at SemEval-2026 Task 11: Neuro-Symbolic Content Stripping for Debaised Syllogistic Reasoning. In Proceedings of the 20th International Workshop on Semantic Evaluation (2026) (pp. 3408–3421). Association for Computational Linguistics.

High-Quality Text-to-Speech Synthesis for the Romanian Language – Mircea Giurgiu

Development of a hybrid method for the automatic recognition of emotions from speech by combining 1D acoustic features and 2D spectro-temporal representations. A 1D–2D CRNN ensemble with late fusion was proposed, complemented by data augmentation strategies to improve robustness to noise and variations in recording conditions. The method was evaluated on the RAVDESS and SAVEE speech datasets, including cross-corpus configurations, highlighting the potential of combining complementary representations to mitigate the effects of domain shift. [1]

Development of an assistive application for text and speech generation through exclusive control of the interface using eye gaze. The system integrates robust gaze direction estimation, signal filtering, and temporal user-intent validation mechanisms, reducing accidental selections. The bilingual Romanian–English interface includes navigation through thematic pages, a virtual keyboard with word prediction, customizable phrases, and speech synthesis, providing an adaptable solution for assistive communication. [2]




Development and comparative evaluation of modern text-to-speech synthesis architectures (Glow-TTS, VITS, FastSpeech2, and Toucan), trained on the Romanian MARA and SWARA corpora. An automated data preprocessing pipeline and a software framework were developed for the objective and reproducible evaluation of the quality, intelligibility, and prosody of synthesized speech. The experimental results demonstrated high model performance, with VITS standing out in terms of spectral fidelity and perceptual quality, while Glow-TTS achieved notable performance in prosody reconstruction.

[1] Valentin-Dimitrie Popescu, Mircea Giurgiu. Hybrid 1D-2D Deep CRNN Ensemble with Feature Selection for Cross-Corpus Robust Speech Emotion Recognition, in *IEEE International Conference on Automation, Quality and Testing, Robotics (AQTR)*, 21-23 Mai 2026, Băile Felix, Romania.


[2] Maeva Cărbunescu, Mircea Giurgiu. AI-driven Gaze Tracking Framework for Text and Speech Generation in Assistive Communication, in *IEEE International Conference on Automation, Quality and Testing, Robotics (AQTR)*, 21-23 Mai 2026, Băile Felix, Romania.

Artificial Intelligence in Robotic Systems – Doina Pîslă


Development of a **dynamic navigation system for dental implantology**, based on markerless detection of anatomical landmarks and their alignment with preoperative 3D models. The implementation of dental segmentation using YOLOv8-seg enables the generation of stable and accurate masks, with a significant reduction in latency compared to the YOLO+SAM approach. The integration of detection, segmentation, 3D visualization, and robotic control into a unified workflow enables the system to be used in real-time surgical navigation scenarios. [1][2]





Development and experimental validation of a **compact spherical robot with a remote center of motion (RCM)**, designed for the orientation of endoscopic cameras and surgical instruments in minimally invasive surgery. A prototype incorporating an RCM mechanism driven by a differential gearbox with counter-rotating motion was developed and integrated, together with embedded control, teleoperated operation, and endoscopic video feedback. Evaluation on a phantom thorax model and using optical tracking confirmed the integrated operation of the subsystems and the feasibility of the proposed solution at the proof-of-concept level. [3]



Development of an **intelligent hand rehabilitation system** based on a wearable exoskeleton and Edge AI, designed for therapy monitoring and adaptation. A hybrid ResMLP–FractalNN architecture was proposed for the automatic classification of hand states based on muscle signals, complemented by XAI methods (SHAP and LIME) for interpreting the model's decisions. The integration of local processing on the NVIDIA Jetson platform lays the groundwork for an adaptive system capable of detecting abnormal behaviors and providing objective and interpretable information on rehabilitation progress.



Development and experimental validation of an experimental **platform for pancreatic robotic surgery** by integrating the ATHENA parallel robot with a flexible laparoscopic instrument, teleoperated control, and a Digital Twin framework based on virtual reality. The integration of an artificial intelligence module enables trajectory refinement and motion assistance in a *human-in-the-loop* setting. Experimental validation on 3D anatomical phantoms demonstrated the execution of maneuvers relevant to minimally invasive pancreatic surgery, including the manipulation and retraction of anatomical structures. [4]




- [1] Paul Tucan, Mihaela Hedesiu, Vasile Bulbucan, Rares Mocan, Daria Pisla, Calin Vaida, Bogdan Gherman, Cristian Dinu. AI-Based Anatomical Landmark Recognition for Robot-Assisted Dynamic Dental Implant Surgery, in *2025 IEEE 21st International Conference on Intelligent Computer Communication and Processing (ICCP)*, Cluj-Napoca, Romania, 2025.
- [2] Bulbucan, V.; Pisla, D.; Tucan, P.; Dinu, C.; Vaida, C.; Mocan, R.; Baciut, M.; Stoia, S.; Hedesiu, M.; Zima, I.; et al. Real-Time Markerless Tooth Detection Towards Dynamic Robot-Assisted Dental Implant Navigation. *Dent. J.* 2026, 14, 345. <https://doi.org/10.3390/dj14060345>
- [3] C. Vaida, D. Horvath; I. Zima, M. Miclaus; B. Gherman, C. Radu, P. Tucan, S. Vegh, D. Sebeni, A. Pisla, et al. Modeling and Design of a Spherical Remote Center-of-Motion Surgical Robot. *Technologies* 2026, 14, 440. <https://doi.org/10.3390/technologies14070440>
- [4] Covaciu, F.; Hajjar, N.A.; Iordan, A.-E.; Corina, R.; Gherman, B.; Cailean, A.; Ciocan, A.; Pusca, A.; Tucan, P.; Pisla, D. Digital Twin-Based Virtual Reality Framework for Interaction and AI-Assisted Control of a Parallel Surgical Robot. *Sensors* 2026, 26, 4410. <https://doi.org/10.3390/s26144410>

AI-based robotics and nonlinear control – Lucian Buşoni


The HRIA activities in **AI-based robotics and nonlinear control** focus on two main directions: resilient control of networked robots and AI-enabled perception/navigation for mobile robotic platforms. Research includes cybersecurity for robotic systems, data-driven nonlinear and optimal control, biomechanical state estimation, and autonomous navigation, together with the development of a multimodal field-robot platform for agricultural applications.

A key result is a **resilient control strategy for drones under cyberattacks**, addressing denial-of-service and man-in-the-middle attacks through a joint anti-windup and saturation-limit design with formal

stability guarantees. The method was validated both in simulation and in real flight, reducing actuator saturation time and maintaining the drone within a safe operating region.



Another major contribution is a **direction-aware object detector for aerial imagery**, extending YOLOv8 with an additional direction-estimation branch. The solution achieves state-of-the-art performance with very limited model-size increase, generalizes to harder datasets without retraining, and was successfully applied to vineyard leaf analysis for future disease-detection and dataset-generation tasks.




The work resulted in several 2025–2026 publications, industrial collaboration with **Envipco**, dissemination through an HRIA special session at IEEE AQTR 2026, and multiple new European and national project proposals.

The contributions combine robust AI-based control, robotic perception, autonomous navigation, and agricultural robotics, with both scientific validation and growing industrial and European impact.

- [1] A. Codrean, O. Ștefan, Z. Lendek, Networked control with anti-windup based resilience to cyberattacks, *International Journal of Systems Science*, 2026
- [2] S. Molnar, L. Tamas, Z. Kato, Directional Object Detection in Aerial Images, *IEEE Geoscience and Remote Sensing Letters*, vol. 23, 2026, art. 6012005 (presented at MIGARS 2025, Bucharest)
- [3] M. Makem, L. Tamas, L. Bușoniu, A reliable approach for identifying acute lymphoblastic leukemia in microscopic imaging, *Frontiers in Artificial Intelligence*, vol. 8, 2025
- [4] Z. Nagy, C. Morărescu, L. Bușoniu, Consensus Value Approximation for a Class of Continuous-Time Opinion Dynamics Models with Stubbornness, *European Control Conference*, 2026

Applications for Satellite Imagery – Dorian Gorgan

Development of automated methods for **soil and agricultural land analysis** based on Sentinel-2 satellite imagery and geospatial and climatic information. Datasets were created and extended by integrating information on soil moisture, precipitation, temperature, elevation, slope, and terrain aspect, together with spectral indices such as NDVI, BSI, NDBI, EVI, SAVI, and NBR. Preprocessing and completion of the time series enabled the generation of consistent datasets for training machine learning models, such as Random Forest, SVM, and Gradient Boosting, for the automated classification of soil and monitoring of agricultural land [1].



[1] Raul-Alexandru Gorgan, Dorian Gorgan, Adaptive Neural Networks for Remote Sensing Imagery: A Systematic Review, *Remote Sensing*, July, 2026

AI-assisted materials design - Adrian Groza

The work introduces **AlloyGraph**, an open-source AI platform for **nickel-based superalloy design and analysis** that integrates an OWL ontology and knowledge graph, physics-informed machine learning, and a multi-agent LLM architecture. The system supports mechanical-property prediction, inverse alloy design, and natural-language querying, while grounding AI reasoning in experimental and metallurgical knowledge.

A rigorous evaluation on **88 commercial superalloys** revealed hidden compositional duplicates under different trade names and therefore introduced an evaluation methodology based on **compositional novelty**. Knowledge-graph anchoring reduced yield-strength error by about **30% for near-duplicate alloys**, while multi-agent reasoning provided its largest benefit for genuinely novel compositions, improving yield-strength prediction by about **18%**. Deterministic physics-based corrections proved most effective for elastic modulus.

The knowledge-grounded research assistant achieved **91.2% accuracy on 250 domain questions**, compared with approximately 50% for ungrounded LLMs. The study also reports important negative findings: agent reasoning can degrade elastic-modulus predictions, and the added value of the LLM Designer for inverse design was not demonstrated against a strong deterministic optimization baseline [1].

AlloyGraph demonstrates a reproducible approach to trustworthy AI-assisted materials design, combining knowledge graphs, physics, machine learning, and multi-agent reasoning, while emphasizing rigorous evaluation, uncertainty estimation, provenance verification, and transparent reporting of both positive and negative results.

[1] Alexandru Lecu, Lezan Hawizy, Soran Biroasca, Adrian Groza, AlloyGraph: A Knowledge-Graph-Grounded Multi-Agent System for Nickel-Based Superalloy Design, <https://github.com/airi-utcn/AlloyGraph>

AI-assisted methodology for the design and optimization of analog and mixed-signal integrated circuits – Marius Neag

The work proposes an **AI-assisted methodology for the design and optimization of analog and mixed-signal integrated circuits**, combining designer expertise with automated optimization tools. Unlike purely “black-box” approaches, the method introduces additional designer-defined constraints and internal circuit requirements to guide the optimization process, making circuit sizing more efficient and producing more robust solutions. The methodology is implemented using industry-standard tools such as **Cadence Maestro EDA**, combining global and local optimization algorithms.

The approach was applied to the design of **multi-GHz frequency dividers for automotive applications** within the EuroDrives project, in cooperation with Robert Bosch, and was further extended to a semi-automated flow for technology migration and circuit redesign.


A representative result concerns the redesign of an automotive **low-dropout voltage regulator (LDO)**. Using BFGS and Conjugate Gradient optimization, the maximum output current was doubled from **100 mA to 200 mA** without increasing quiescent current, while also achieving better transient response than a conventionally designed solution.

The contribution demonstrates that combining expert knowledge with AI-based optimization can significantly improve the efficiency, robustness, and performance of analog IC design and redesign workflows.

[1] Paul Coste, Paul Miresan, Marius Neag, Cristian Raducan, Marina Topa, Istvan Kovacs – "Automotive Capacitor-Less Linear Regulator with Active Capacitance Multiplier Frequency Compensation and AI-Assisted Sizing" – *Romanian journal of information science and technology (ROMJIST)*, Volume 29, Number 2, 2026, pp. 143–154.

Privacy-Preserving Radar–Audio Fusion for Indoor Intrusion Detection / Edge AI Prototype for Radar–Audio Indoor Intrusion Detection

A privacy-preserving multimodal sensing prototype was developed for indoor intrusion detection using complementary 60 GHz radar and audio information. A dedicated radar–audio dataset was collected in a controlled indoor environment, covering intrusion-related events together with challenging non-intrusive scenarios such as ambient noise, external sirens, speech, falling objects, and other potential false-alarm sources. The resulting dataset, Radar-Audio Intrusion Dataset for Privacy-Preserving Indoor Alarm Detection, was publicly released through Zenodo to support reproducible research.



Several AI processing strategies were evaluated, including radar-only and audio-only detection, feature-level early fusion, and score-level late fusion. The strongest window-level result was obtained using late radar–audio fusion, reaching an F1-score of 0.897. A stateful alarm mechanism based on temporal smoothing, hysteresis, and persistence reduced false-positive windows from 4.19 to 2.58 per 100 negative windows while retaining an F1-score of 0.888.

A functional edge prototype was implemented and evaluated on a Raspberry Pi 5, reproducing the reference test decisions without platform-specific retraining. The selected late-fusion classifier payload requires approximately 1.16 MB, while stabilized raw-audio-to-decision replay with pre-windowed radar input requires approximately 1.30 ms per window. These results demonstrate the feasibility of compact edge-


oriented radar–audio intrusion detection. The work is currently being prepared for publication in an international scientific journal.

[1] E. Vladu, L. Rostas, A. Battigelli, D. Domnici, M. Neag, Radar-Audio Intrusion Dataset for Privacy-Preserving Indoor Alarm Detection, Zenodo, Version 1.0, 2026, <https://doi.org/10.5281/zenodo.21885389>


Digital Circuits and Systems Design – Sorin Hintea

The HRIA activities carried out by the **Digital Circuits and Systems Design team** focus on three main directions: AI-enabled cardiac monitoring, multimodal biometric recognition, and intelligent self-interference cancellation for full-duplex wireless communications.

For **cardiac monitoring**, an embedded/FPGA-oriented ECG analysis chain was developed, covering signal-quality assessment, artifact detection, and arrhythmia classification. Random Forest achieved up to **98.18% accuracy** for artifact detection, while the Xception–Amor combination reached **99.3% accuracy** for eight rhythm classes. The work also investigated FPGA implementation strategies based on CNNs, quantization, pipelining, pruning, and dataflow to support low-latency, energy-efficient local inference.



A second direction concerns **multimodal biometric recognition using ECG, PPG, and PCG signals**. A bimodal ECG–PPG system exceeded **90% identification accuracy**, while additional work addressed automatic PCG-based heart-murmur detection and FPGA-based AES-128 protection of physiological data, creating a basis for compact and secure biometric systems.



The third direction focuses on **In-Band Full-Duplex wireless communications**, for which an experimental hardware platform was developed using two Xilinx ZC702 FPGA boards and AD-FMCOMMS3 RF modules. The demonstrator supports simultaneous transmission and reception in the same frequency band and enables the evaluation of self-interference reduction techniques.

The work demonstrates strong integration of AI, embedded/FPGA implementation, biomedical signal processing, cybersecurity, and advanced wireless communication, supported by scientific publications and a functional hardware demonstrator.

Results – Datasets, Technologies, Demonstrators, Prototypes

1. Datasets

<i>ID</i>	<i>Dataset</i>	<i>Link</i>	<i>QR Code</i>
Environment Perception, Representation and Understanding from Aerial Perspective – Sergiu Nedevschi and Florin Oniga			
1	Synthetic Aerial Dataset - ClaraVid	https://rdbch.github.io/claravid	
2	Aerial 3D Reconstruction Dataset - UAVid-3D-Scenes	https://github.com/hrflr/uavid-3d-scenes	
3	Aerial Semantic Segmentation Dataset - UAVid++	https://vivichiciudean.github.io/uavidplusplus/	
4	Synthetic Aerial Dataset - Forest Inspection Dataset (Sunny sequences)	https://zenodo.org/records/15511426	
5	Synthetic Aerial Dataset - Forest Inspection Dataset (Overcast Sequences)	https://zenodo.org/records/17912570	

2. Technologies

<i>ID</i>	<i>Technology</i>	<i>Link</i>	<i>QR Code</i>
Environment Perception, Representation and Understanding from Aerial Perspective – Sergiu Nedevschi and Florin Oniga			
1	Delentropic Scene Profile	https://docs.google.com/document/d/1AiQ0a_KDWWfle_wF1GyH8oQn10NCZVNzoyrxu5u1KcE/edit?usp=sharing	

2	Image Complexity in Neural Scene Reconstruction	https://docs.google.com/document/d/1Ae7letHHhV-T35nTGr8luuyrWku7C9Xgvra5ers2-4k/edit?usp=sharing	
3	Terrain-Conditioned Metric Depth Estimation for UAVs	https://didatec-my.sharepoint.com/:w:/g/personal/fl_orea_r_horatiu_campus_utcluj_ro/IQBD2f0ngQgiS4dJv4Z5XV0kAe5C-DfN3aDEIpfYt0rPjLs?e=pAg2mp	
4	Semantic Segmentation for Aerial Imagery - DINO-UNet	https://docs.google.com/document/d/1DZOc4GZ8n8lDa3jBc-byU9PEzR7AAjD7LALXecCFJ0U/edit?usp=sharing	
5	Aerial Training and Inference Strategy	https://docs.google.com/document/d/1bcsiA8jUF9xyG1JinukJSVDreZmcZHodg6AqCCw_QHY/edit?usp=sharing	
6	Change Detection in Aerial Images	https://docs.google.com/document/d/16T-sKt91tN-l5P2m5pM-NUDR0uH0kIXARQlkgThtHHE/edit?usp=sharing	
7	Visual Question Answering in Aerial Images	https://docs.google.com/document/d/1g32ADn0OPQFNU6LqlPEZ1HVxe9PrXSajUJqUDWEgAJY/edit?usp=sharing	
Medical Image processing – Tiberiu Marița			
8	Method for classifying diabetic retinopathy from retinal Angio-OCT images	https://docs.google.com/document/d/1ENeLBRxn6O9Xy8dBxxa9JYDnLgLd8yQJLZdVACb9VPU/edit?usp=drive_link	

9	Retinal vascular network segmentation method in OCTA images using a SAM-based model	https://docs.google.com/document/d/1UO2802kK_WD5adrdioJT-WaAD04cbOX4Q8TO6kUIRnI/edit?usp=drive_link	
10	Vertebra segmentation method for X-ray images using a SAM-based model.docx	https://docs.google.com/document/d/1HIOKRtNdLyzFS56GVTkeS7k8W6IzqGC1ZaOwIvLXOJ8/edit?usp=drive_link	
11	Method for detecting thoracic pathologies in X-ray images	https://docs.google.com/document/d/15vqr0zJjXU5Zg0RID6IjqcSZHk0uq19jz5huknIQ-8/edit?usp=drive_link	
Intelligent Applications for Smart Cities and the Environment – Cioară Tudor			
12	Representation-Augmented AI Model for Energy Forecasting	https://dsrl.eu/hria/ai-technologies/augmentation/	
13	Physics-Constrained AI Imputation	https://dsrl.eu/hria/ai-technologies/imputation/	
14	LLM-Based Knowledge Graph Extraction from Cloud Logs	https://dsrl.eu/hria/ai-technologies/knowledge-graph/	
Artificial Intelligence in Speech Processing – Mircea Giurgiu			
15	Attention-Driven Speech Emotion Recognition	https://drive.google.com/file/d/1QC_C0HwTZU5J-7Lm5oR8aVJyGVMgtE3RZ/view?usp=sharing	

16	Audio and Text Conversational Speech Emotion Recognition	https://drive.google.com/file/d/1-Qq-XuN40ymLnJk5_1OVcpoO5PLdbMKs/view?usp=sharing	
17	Audio deepfake detection	https://drive.google.com/file/d/1QH-SVqb-H4qX76HXgJ3WPvjMbKfYz82cW/view?usp=sharing	
18	Audiofakes in the Noise	https://drive.google.com/file/d/1R_58ECIxmMjQqw1Gjn-r6S-b8kEhUas/view?usp=sharing	
19	Automatic Objective Evaluation of TTS	https://drive.google.com/file/d/1digs7JsUsg_ku_JqjA37knQL_M1Ik0p/view?usp=sharing	
20	Gaze2Speech	https://drive.google.com/file/d/1Q3EuxjdNWTzID_znDdj7W4E4xhgKGo65/view?usp=sharing	
21	Hardware Platform for UWB- Based Spatial Audio Rendering	https://drive.google.com/file/d/11EnDCxcHc_WcfAFLDw33wrBCjpPmAwbW/view?usp=sharing	
22	Hybrid CRNN for Cross-Corpus Emotion Recognition	https://drive.google.com/file/d/1bzBV7nXjUdqEtEYVMoUTvU3wKrusBNUR/view?usp=sharing	

23	Low-Power Embedded Speech Enhancement	https://drive.google.com/file/d/1J4uVWPmROrGQ9e7BEz8SS6Gerb-LWZ3e/view?usp=sharing	
24	Objective metrics for evaluation of Romanian TTS systems	https://drive.google.com/file/d/1ElwANFE5oXdLJmLxsjDYRafaZSjJnotq/view?usp=sharing	
25	Speaker Anonymization via GAN-driven embeddings	https://drive.google.com/file/d/1TbWCQf190acX4lSDIC0kzJUuDKoimPj5/view?usp=sharing	
26	Speech and Text Watermarking	https://drive.google.com/file/d/116XlJv9z1nq-bMjXVT5Wkis0_TzJjHKi/view?usp=sharing	
27	Text and Speech Generation in AAC	https://drive.google.com/file/d/10YWHpHVhSFfN_WKAnSbzRElSqwD3r21C/view?usp=sharing	

3. Demonstrators

ID	Demonstrator	Link	QR Code
AI-based robotics and nonlinear control – Lucian Buşoniu			
1	Vision-based Navigation and Action Models: Benchmarking the Existing Approaches	https://docs.google.com/document/d/1IIVi7-ZmcztmrCNUrwHhxxu3o0dG9ux5sndAX7VbtHnM/edit?tab=t.0	

Medical Image processing – Tiberiu Marița			
2	Demonstrator prototype for segmentation and classifying of anatomical structures in multimodal medical images	https://docs.google.com/document/d/1-t1-1uptjrlv-32nOCvtaLp0Bvr6vaKtEfNYV2vn78w/edit?usp=drive_link	
Intelligent Applications for Smart Cities and the Environment – Cioară Tudor			
3	AI-Based Energy Forecasting Demonstrator	https://dsrl.eu/hria/energy-prediction-demo	

4. Prototypes

ID	Prototype	Link	QR Code
AI-assisted methodology for the design and optimization of analog and mixed-signal integrated circuits – Marius Neag			
1	Edge AI Prototype for Radar–Audio Indoor Intrusion Detection	https://doi.org/10.5281/zenodo.21885389	
AI-assisted materials design - Adrian Groza			
2	A Knowledge-Graph-Grounded Multi-Agent System for Nickel-Based Superalloy Design	https://github.com/airi-utcn/AlloyGraph	