EURA NOVA

INTERNSHIP OFFERS 2027

FACULTIES OF SCIENCE AND ENGINEERING

EURANOVA

Welcome to Euranova's Internship Program

Euranova is an international consulting company that specialises in data-driven business solutions, guided by a strong company culture and a passion for innovation. Founded in September 2008 and located in Brussels, Marseille, and Tunis, our purpose is to bring life to our customers' great ideas by offering best-in-class services in data science, engineering, and architecture.

Inside and out, we care for collaboration, impact, and excellence, and in line with this course of action, we offer academic programs in partnership with universities. See below for details regarding our internship offers.

Explore Our Internship Opportunities

This document presents internship topics supervised by our consultancy and innovation department. Each project is an opportunity to be actively involved in the development of solutions to address tomorrow's challenges in ICTs and to implement them today, and can be developed into end-of-study projects or master's theses. The students will work in a dedicated international team of engineers with diverse expertise in machine learning, computer vision, Large Language Models (LLMs), data privacy, AI governance, high-performance computing, etc.

Our company culture is built on mutual enrichment, continuous learning, empathy in coding, and knowledge sharing. Though we consider that attitudes take precedence over technical prowess, we usually work with solution-oriented engineers in computeré science and applied sciences, who thrive in a collaborative environment and embrace challenges with humility, curiosity, and a scientific mindset. To learn more about our R&D activities or company culture, visit our website at https://euranova.eu.

How to apply

Interested in being a part of our story? Here's how to move forward:

- When you have gone through our internship offers, pick your favourite three.
- Draft a short text for each one, stating why you find it interesting and what you would do about it.
- Send us this statement, along with your CV, to career@euranova.eu.

Please note that the locations and dates are indicative; do not hesitate to contact us to find an arrangement. Although previous experience with the technologies mentioned in the offers is appreciated, it won't be amongst the main criteria for intern selection.

Submission deadline: we encourage you to apply early, as we will review applications as they come in to find the best matches. Should you have any questions regarding our internship offers or the selection process, do not hesitate to reach out to career@euranova.eu.

Co-Generation And Learning To Solve Constraint Programming Problems

1. Context

Although many real-world problems may be simple to state with natural language, finding a solution to these problems automatically may be difficult. Constraint Programming (CP) has been for tens of years now the proper way to formalize a problem and its constraints in a dedicated technical language, which then allows some black box solvers to take over and automatically find optimal solutions for us. However, formalizing a problem into a constrained program can be tricky, which is the reason why the research community has looked for ways to automate going from a natural language description of a problem to a proper formalization. Large Language Models (LLMs) are a promising step in this direction.

An internship during summer 2026 showed it is possible to make LLMs quite good at modeling relatively simple constrained programming problems without any training/fine-tuning, by framing them in a proper harness. In particular, LLMs that are initially bad at producing such programs became able to do so successfully when enclosed in the right system. However, this activity kept hitting a wall again and again: proper datasets of natural language descriptions paired with an associated reference constrained program are rare on the internet, which made us create our own during the summer 2026 internship. By the date of this writing, this benchmark is considered beaten (100% success using cheap LLMs within a proper harness to do the task of formalizing CP problems). However, it does not mean we have cracked open the problem of formalizing CP problems, only that we don't have difficult-enough problems in this benchmark to quantify how good our approaches have become.

The goal of this 2027 internship is to design methodologies to co-design a new dataset of {natural language problems, associated CP program} and optimized harnesses that can crack industry-level optimization problems. In order to do so:

- The intern will first design theoretical and empirical ways to quantify the difficulty of modeling a specific natural language problem. Amongst other possibilities, all the data generated during the 2026 internship (natural language problems, reference CP models, generated CP models, additional metadata) will be usable to that end.
- Based on these definitions of difficulty, we will adapt traditional curriculum-based optimizations (solving harder and harder problems) to the natural language world, to adversarially generate more and more difficult problems, and adapt more and more sophisticated harnesses to solve them.

The ultimate goal of the internship will be to crack open anonymised versions of optimization problems Euranova has solved in the past in the industry.

2. Technologies

- Constraint programming: Although the topic can be handled with any solver, we will focus on CPMpy (no prior knowledge required on this framework)
- Large Language Models (LLMs): OpenRouter, OpenAI API, LLM Harness
- Natural language program optimization: DSPy, GEPA, "System 1" models

3. Objectives

- **Literature review:** In the fast-moving world of LLMs, a study of approaches available at the time of the internship will be studied. Those insights will be the grounding of new ideas to experiment with.
- **Design of evaluation methodologies:** Evaluating what we are doing, and the progress between iterations, is paramount, and can be quite tricky for this kind of problem. The intern will have to find ways to quantify progress.
- **Design and implementation of solutions:** Starting from the currently available internal repository, the intern will design and implement approaches to generate more and more difficult problems, and associated modeling harnesses.
- **Experimental Benchmarking:** Based on the designed evaluation setup, the intern is expected to provide insights on what works, what does not, and what impact each iteration on the project brings.

4. Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2027.

5. References

[1] Michailidis, K., Tsouros, D. and Guns, T. "CP-Bench". Zenodo. Available at: https://doi.org/10.5281/zenodo.15592407, 2025.

[2] Singirikonda, Akash, Serdar Kadioğlu and Karthik Uppuluri. "Text2Zinc: A Cross-Domain Dataset for Modeling Optimization and Satisfaction Problems in MiniZinc." *arxiv preprint at* abs/2503.10642, 2025.

[3] Xianchao Xiu and Jianhao Li and Jun Fan and Wanquan Liu, "Large Language Models for Operations Research: A Comprehensive Survey", *arxiv preprint at* https://arxiv.org/abs/2605.20849, 2026.

[4] Yuliang Song and Eldan Cohen, "Do LLMs Understand Constraint Programming? Zero-Shot Constraint Programming Model Generation Using LLMs", *THE 19TH LEARNING AND INTELLIGENT OPTIMIZATION CONFERENCE*, https://openreview.net/forum?id=6zlpzSKzqj, 2025.

Multi-Frame Temporal Convolution for Real-Time Object Detection

1. Context

Standard real-time detectors (e.g., YOLO) process video frame by frame, missing temporal cues essential for identifying small targets like drones. Conversely, spatio-temporal attention architectures are too computationally heavy for edge NPUs like the NXP i.MX8M Plus (~30 ms target latency). While recent baselines such as Quan et al. (2025) introduce lightweight multi-frame integration for YOLO, they do not address edge-hardware operator barriers or sensor platform dynamics.

This internship addresses these gaps with two core contributions:

- NPU-Native Temporal Fusion: Designing 2D temporal fusion primitives strictly compatible with accelerated NXP eIQ/TFLite operators, eliminating CPU fallback and memory-shuffling overhead.
- Ego-Motion Decoupling: Learning temporal representations capable of distinguishing physical drone motion from background motion induced by moving camera platforms.

2. Technologies

- Architectures & Paradigms: Lightweight 2D Detectors (e.g. YOLOX), Temporal Shift Module (TSM), Multi-frame Early/Feature Fusion
- Edge Hardware & Deployment: NXP i.MX8M Plus NPU, NXP eIQ Toolchain, TensorFlow Lite, ONNX
- Frameworks & Libraries: Python, PyTorch, OpenCV
- Datasets: USC Drone, Anti-UAV RGBT
- Mining & Annotations Pipeline: Single-frame reference teachers (e.g. RT-DETR) with human verification tooling (e.g. CVAT)

3. Objectives

- Literature & Operator Audit: Screen lightweight temporal methods (TSM, frame differencing, StreamYOLO) against the NXP eIQ operator library to verify native NPU acceleration and eliminate CPU fallback.
- Dataset & Benchmark Curation: Curate a multi-frame drone benchmark (Anti-UAV RGBT, USC Drone) and build a failure-mode test set via automated RT-DETR candidate mining paired with human ground-truth validation.
- Architecture Design: Integrate lightweight temporal shift and fusion primitives (building on Quan et al. and TSM) into a 2D detector (e.g., YOLOX), optimized for NPU kernel constraints and resilient to camera ego-motion.

- Quantization & Edge Optimization: Export the model to TFLite and deploy via NXP eIQ, applying INT8 post-training quantization (PTQ) and quantization-aware training (QAT) on the i.MX8M Plus NPU.
- Experimental Benchmarking: Evaluate trade-offs between frame-level mAP, temporal detection rate (P_d), false alarms per hour, and latency (FPS/ms) against single-frame baselines across fixed and moving camera setups.

4. Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2027.

5. References

[1] J. Lin et al., "TSM: Temporal Shift Module for Efficient Video Understanding," Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 7083-7093, 2019.

[2] J. Redmon et al., "You Only Look Once: Unified Real-Time Object Detection," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 779-788, 2016,

[3] J. Yang et al., "StreamYOLO: Real-time Object Detection for Streaming Perception," IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2022.

[4] Z. Ge et al., "YOLOX: Exceeding YOLO Series in 2021," arXiv preprint arXiv:2107.08430, 2021,

[5] Y. Quan et al., "Lightweight Multi-Frame Integration for Robust YOLO Object Detection in Videos," IEEE European Conference on Mobile Robots, 2025,

[6] Y. Zhao et al., "DETRs Beat YOLOs on Real-time Object Detection," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 16965–16974, 2024.

Vendor-Agnostic AI Harness for Agentic Data Science Workflows

1. Context

Modern AI coding agents struggle with the non-linear, state-heavy nature of Data Science when restricted to standard linear context windows, leading to context rot, token bloat, and execution loops. This research internship investigates state-management architectures tailored for agentic Data Science workflows.

To prevent vendor lock-in, the project decouples orchestration from LLM providers by building a vendor-agnostic harness. It connects open-source orchestrators (e.g. OpenCode, Hermes, Pi) to enterprise-compliant LLM endpoints (e.g. Mistral AI API, Vertex AI EU) paired with local sandboxed code execution.

The scope is bounded to a 2-pass iterative loop:

- Pass 0 (Clean Slate): EDA → Feature Engineering → Baseline Model.
- Pass 1 (Refinement): Modifying the generated codebase based on holdout error feedback.

The intern will conduct empirical ablations comparing three state-management paradigms:

- Linear Prompt Rewind: Rolling back history upon execution failure.
- Workspace Snapshots: Restoring full workspace file-system states before refactoring.
- Event-Replay State Trees: Maintaining branching histories for dynamic backtracking.

2. Technologies

- Orchestration & State: Open-source agent frameworks (e.g. OpenCode, Hermes, Pi), Event-driven State Trees, LangGraph
- LLMs & Enterprise APIs: Mistral AI API (EU Sovereign endpoints), Google Vertex AI (Gemini EU endpoints)
- Sandboxing & Execution: Docker / WASM local sandboxing, Jupyter Kernel execution APIs
- Data Science: Python, Scikit-Learn, LightGBM, PyTorch, Pandas

3. Objectives

- Literature Review: Audit state-of-the-art non-linear state management and context pruning techniques in coding agents.
- Benchmark Curation: Assemble an evaluation suite from post-cutoff ML challenges (e.g. Kaggle) to prevent LLM data contamination.
- Harness Development: Build a decoupled harness implementing sandboxed local execution and the three state-management paradigms.

- Experimental Benchmarking: Quantify trade-offs across downstream model quality (M), token efficiency (N_{tokens}), autonomous error recovery (R_{auto}), and network egress compliance ($L_{compliance}$). As a secondary evaluation, assess human steering overhead (F_{human}) through a qualitative pilot audit with the supervisors and intern.

4. Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2027.

5. References

[1] Noah Shinn et al., "Reflexion: Language Agents with Verbal Reinforcement Learning", Advances in Neural Information Processing Systems 36 (NeurIPS 2023), https://proceedings.neurips.cc/paper_files/paper/2023/hash/1b44b878bb782e6954cd888628510e90-Abstract-Conference.html.

[2] John Yang et al., "InterCode: Standardizing and Benchmarking Interactive Coding with Execution Feedback", Advances in Neural Information Processing Systems 36 (NeurIPS 2023 Datasets and Benchmarks Track), https://www.proceedings.com/075280-1035.html.

[3] Jun Shern Chan et al., "MLE-bench: Evaluating Machine Learning Agents on Machine Learning Engineering", The Thirteenth International Conference on Learning Representations (ICLR 2025), https://openreview.net/forum?id=6s5uXNWGIh.

[4] Sirui Hong et al., "Data Interpreter: An LLM Agent For Data Science", Findings of the Association for Computational Linguistics: ACL 2025, https://aclanthology.org/2025.findings-acl.1016/.

[5] Zhengyao Jiang et al., "AIDE: AI-Driven Exploration in the Space of Code", arXiv preprint at https://arxiv.org/abs/2502.13138, 2025.

New Debiasing Methods for Combating Acquisition Bias in AI for Medical Image Segmentation

1. Context

Deep learning models—particularly fully convolutional networks and vision transformers (e.g., U-Net, nnU-Net, Swin UNETR)—have achieved state-of-the-art performance in medical image segmentation, more specifically in tumor delineation. However, these algorithms are highly susceptible to taking shortcuts by learning unwanted representations from sub-population shifts, imaging hardware variability, and demographic attributes such as sex, age, ethnicity, acquisition protocol, or scanner vendor.

Consequently, when deployed across heterogeneous clinical settings, models often exhibit severe performance disparities across patient cohorts, leading to unfair diagnostic outcomes and unreliable segmentation boundaries for underrepresented groups. Furthermore, retraining complex segmentation models from scratch whenever bias or unwanted samples are detected is mostly impossible due to the lack of segmentation labels.

To address this issue, recent research explores fairness in deep medical image processing:

- Adversarial debiasing and disentanglement: Early domain-unlearning approaches rely on adversarial training (for instance, gradient reversal layers) to force feature extractors to become invariant to dataset/scanner membership, effectively "unlearning" domain bias during representations learning.[1]
- Fairness-constrained segmentation losses: Augmenting standard dense prediction losses (Dice, Cross-Entropy) with group-fairness penalties (e.g., demographic parity, equalized odds mapped to pixel-level metrics) or invariant risk minimization objectives.[2]
- Distribution-aware architectural modules: Employing adaptive gating or mixture-of-experts modules to explicitly model data heterogeneity and decouple clinical segmentation features from demographic shifts[3] .

Despite recent advances, unlearning bias specifically in segmentation tasks of tumor delineation remains under-explored compared to standard image classification. Indeed, because bias creates patchy, localized errors across different parts of the image, segmentation tasks are much harder to fix than simple labeling tasks. Moreover, combating biases without damaging the fine-grained boundary segmentation capability and overall model generalization remains a fragile trade-off.

The goal of this internship is to formalize, benchmark, and propose a novel bias unlearning method tailored to state-of-the-art medical image segmentation models. The proposed approach will be applied on tumor segmentation within 3D medical images.

[1] Kamnitsas, K., et al. (2017). Unsupervised domain adaptation in brain lesion segmentation with adversarial networks. International Conference on Information Processing in Medical Imaging (IPMI).
[2] Tian, Y., et al. (2024). Harvard FairSeg: A Large-Scale Medical Image Segmentation Dataset for Fairness Learning. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) / arXiv:2311.02189.
[3] Oh, S., et al. (2025). Distribution-aware Fairness Learning in Medical Image Segmentation. International Conference on Machine Learning (ICML).

2. Scope of Work

The internship will have the following objectives, in chronological order:

- Explore state-of-the-art approaches for debiasing methods in medical imaging tumor segmentation.
- Preprocess and explore the data.
- Develop a debiasing method based on the most promising approach.
- Assess its performance, compare it with previously developed methods and discuss the results.

The results of this internship could be presented to the community as a talk, a report or a scientific paper.

3. Skills

- Tools & Version Control: Git/GitLab, Docker
- Programming Languages: Python (Pandas, NumPy,)
- Data Science Libraries: Scikit-learn, PyTorch
- Visualization Tools: Matplotlib, Seaborn, Plotly
- Medical Imaging Knowledge: Basic understanding of medical imaging modalities and AI applications

Although previous experiences using such technologies would be appreciated, this won't be amongst the main drivers of the intern selection.

4. Where and when

France (Marseille), from March/April to August/September 2027.

Open-Vocabulary 3D Instance Segmentation for Industrial Asset Inventory

1. Context

Modern 3D instance segmentation frameworks (such as Mask3D or SoftGroup) rely on fully supervised, closed-set training over fixed category taxonomies. In industrial facility inspection, power utility monitoring, and plant digital twinning, these models fail because unannotated industrial components (valves, pressure vessels, pumps, custom structural brackets) lie in the long tail of visual distributions. Dense point-level 3D manual annotation in such complex, brownfield industrial environments is economically prohibitive and practically unfeasible. Furthermore, standard open-vocabulary feature fields (e.g., LERF, raw CLIP/DINOv2) lack sharp 3D spatial instance boundaries, frequently causing over-segmentation or the accidental merging of adjacent physical components.

To resolve these limitations, this internship aims to develop an end-to-end 3D Gaussian Grouping pipeline that performs class-agnostic 3D instance discovery and open-vocabulary asset inventory without requiring manual 3D labels. By assigning low-dimensional per-Gaussian instance embeddings ($f_i \in \mathbb{R}^d$) supervised by 2D zero-shot masks from foundation models (SAM2) via contrastive prototype-pull, margin-push, and ray-variance regularization losses, the model learns multi-view consistent object boundaries directly from multi-view imagery. Unsupervised clustering (e.g., HDBSCAN) then partitions the Gaussians into discrete 3D object instances, which are grounded with text embeddings (CLIP/VLM codebooks) to enable natural language querying, spatial cataloging, and automated asset inventory.

2. Technologies

- 3D Neural Representations: 3D Gaussian Splatting (3DGS), Gaussian Grouping, 2D Gaussian Splatting (2DGS).
- 2D/3D Foundation Models: Segment Anything Model 2 (SAM2), Grounding DINO, CLIP, OpenSeg, DINOv2.
- Deep Learning & Spatial Frameworks: PyTorch, PyTorch Lightning, gsplat, Open3D.
- Clustering & Embeddings: HDBSCAN, Scikit-learn, FAISS, contrastive loss formulations.
- Benchmarks & Evaluation: ScanNet++, STPLS3D, GauU-Scene V2, industrial plant-room captures.

3. Objectives

- Literature Review & Harness Setup: Review state-of-the-art 3D Gaussian grouping architectures (OpenSplat3D, InstanceGaussian, DCSEG) and establish a standardized benchmark harness evaluating class-agnostic instance AP_{50}, AP_{25}, and mIoU.
- Contrastive Gaussian Grouping Pipeline: Implement per-Gaussian instance embedding channels supervised by SAM2 2D mask projections using prototype attraction, boundary margin push, and ray-variance regularization ($\mathcal{L}_{var}$) to guarantee cross-view identity consistency.

- Class-Agnostic Boundary Discovery: Integrate adaptive clustering algorithms (HDBSCAN / graph connectivity) to discover and segment variable-scale 3D asset instances without relying on pre-set object counts or tracking heuristics.
- Open-Vocabulary Language Grounding: Implement an instance-level language grounding layer (per-instance CLIP pooling and VLM codebooks) enabling natural language prompt querying (e.g., "locate shut-off valves", "highlight power pylon crossbeams") with adaptive thresholding for multi-instance retrieval.
- Multi-View Aerial/Industrial Benchmarking: Validate boundary stability and ID-switch resilience on wide-baseline aerial sequences and evaluate 3D discovery precision on real-world industrial plant-room captures against laser-scanned ground truth.

4. Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2027.

5. References

J. Piekenbrinck, C. Schmidt, A. Hermans, N. Vaskevicius, T. Linder, and B. Leibe, "OpenSplat3D: Open-Vocabulary 3D Instance Segmentation using Gaussian Splatting," IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 5246–5255, Jun. 2025. DOI: https://doi.org/10.1109/cvprw67362.2025.00520

H. Li et al., "InstanceGaussian: Appearance-Semantic Joint Gaussian Representation for 3D Instance-Level Perception," IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 14078–14088, Jun. 2025. DOI: https://doi.org/10.1109/cvpr52734.2025.01314

L. Wiedmann, L. Wiehe, and D. Rozenberszki, "DCSEG: Decoupled 3D Open-Set Segmentation using Gaussian Splatting," IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW), pp. 5217–5226, Jun. 2025. DOI: https://doi.org/10.1109/cvprw67362.2025.00517

H. Pan, B. Liu, X. Xu, C. Zheng, Y. Nie, and S. He, "Gaussian Prompter: Linking 2D Prompts for 3D Gaussian Segmentation," IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), vol. 47, no. 9, pp. 8165–8179, Sep. 2025. DOI: https://doi.org/10.1109/tpami.2025.3576839

M. Kamali, B. Atazadeh, A. Rajabifard, and Y. Chen, "Training-free open-vocabulary monocular 3D object detection for industrial assets," Automation in Construction, vol. 184, p. 106841, Apr. 2026. DOI: https://doi.org/10.1016/j.autcon.2026.106841

GSD-Bounded 3DGS Mesh Extraction for As-Built vs. As-Designed CAD/BIM Auditing

1. Context

Standard 3D Gaussian Splatting (3DGS) achieves exceptional photorealistic view synthesis, but its explicit anisotropic primitives lack continuous topological surface bounds. Because Gaussian ellipsoids are optimized purely against a 2D photometric rendering loss, unconstrained training produces volumetric "floaters" and noisy, unoriented density fields. Consequently, standard Marching Cubes surface extraction yields disconnected triangle soups with substantial geometric reconstruction errors, rendering vanilla 3DGS unsuitable for strict engineering inspections, Scan-to-BIM modeling, or tolerance verification.

In photogrammetry, achievable spatial accuracy is strictly bounded by the sensor's Ground Sampling Distance (GSD, typically 1.5–2×GSD under optimal geometric network constraints). Any claim of "sub-millimeter" precision from drone altitudes is physically unfounded. This internship focuses on building a metrologically sound surface extraction and CAD/BIM compliance auditing pipeline. By replacing volumetric 3D ellipsoids with surface-regularized planar primitives (2DGS / PGSR) constrained by normal-consistency, depth-distortion, and observation-aware TSDF fusion, the pipeline extracts watertight, manifold 3D polygonal meshes with geometric accuracy bounded by physical GSD resolution. The reconstructed as-built meshes are subsequently registered against nominal as-designed CAD/BIM models (IFC 4.3 / Revit) to generate signed 3D surface deviation heatmaps for automated structural tolerance auditing.

2. Technologies

- Surface-Regularized 3DGS: 2D Gaussian Splatting (2DGS), Planar Gaussian Splatting (PGSR), Sparse2DGS, SuGaR.
- Surface Extraction & Geometric Meshing: Truncated Signed Distance Function (TSDF) fusion, Marching Cubes, Adaptive-GradNormal meshing.
- Structure-from-Motion & Photogrammetry: COLMAP, Open3D, PyTorch, CUDA.
- CAD/BIM & Metrology Standards: IFC 4.3 schema, CloudCompare, Iterative Closest Point (ICP), ISO 19650 (CDE information containers), ISO 10360-style tolerance verification.
- Datasets & Benchmarks: GauU-Scene V2 (UAV/LiDAR), Tanks and Temples, DTU, municipal/industrial building scans.

3. Objectives

- Failure-Mode Quantification & Baseline Setup: Benchmark standard 3DGS mesh extraction against planar-regularized baselines (2DGS, PGSR) on open photogrammetric benchmarks, quantifying floater artifact density, surface-normal error, and non-manifold boundary edges.
- GSD-Bounded Surface Reconstruction Pipeline: Implement a surface-aligned 2D/3D Gaussian pipeline incorporating ray-splat intersection, depth-distortion, and normal-consistency losses, guided by monocular normal/depth priors to guarantee planar flatness.

- Observation-Aware Watertight Mesh Extraction: Implement multi-view observation-aware TSDF fusion at a fixed voxel grid bounded by 2×GSD (truncation at 4×GSD), extracting watertight, manifold triangular meshes using Marching Cubes and feature-preserving remeshing.
- CAD/BIM Registration & Deviation Auditing: Develop an automated co-registration workflow (coarse feature alignment + point-to-plane ICP) to register the extracted mesh against nominal CAD/IFC design models, generating signed Euclidean surface deviation heatmaps for building tolerance verification.
- Experimental Benchmarking & Validation: Evaluate mesh accuracy (Chamfer Distance, surface-to-LiDAR RMSE, and F1-score at tolerance thresholds) on UAV benchmarks (GauU-Scene V2) to prove that the reconstructed as-built surface reliably achieves 1.5–2×GSD accuracy.

4. Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2027.

5. References

J. Wu, R. Li, Y. Zhu, R. Guo, J. Sun, and Y. Zhang, "Sparse2DGS: Geometry-Prioritized Gaussian Splatting for Surface Reconstruction from Sparse Views," IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 11307–11316, Jun. 2025. DOI: https://doi.org/10.1109/cvpr52734.2025.01056

B. Huang, Z. Yu, A. Chen, A. Geiger, and S. Gao, "2D Gaussian Splatting for Geometrically Accurate Radiance Fields," in Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pp. 5333–5343, Jun. 2024. DOI: https://doi.org/10.48550/arXiv.2403.17888

H. Xiang et al., "GaussianCraft: Fine-grained 3D Gaussians for efficient large-scene surface reconstruction," ISPRS Journal of Photogrammetry and Remote Sensing, vol. 235, pp. 651–667, May 2026. DOI: https://doi.org/10.1016/j.isprsjprs.2026.03.020

X. Liao, Y. Liu, Y. Huang, Y. Luo, Z. Zhou, and H.-Y. Chong, "Integrating CAD, BIM, immersive technology, and 3D Gaussian Splatting for construction model coordination with ISO 19650," Scientific Reports, vol. 16, no. 1, p. 32244, Dec. 2025. DOI: https://doi.org/10.1038/s41598-025-32244-y

E. B. Anil, P. Tang, B. Akinci, and D. Huber, "Deviation analysis method for the assessment of the quality of the as-is Building Information Models generated from point cloud data," Automation in Construction, vol. 35, pp. 507–516, Nov. 2013. DOI: https://doi.org/10.1016/j.autcon.2013.06.003