



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 computer 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

Current real-time object detection architectures (such as the YOLO family) process video streams frame by frame, ignoring the temporal context crucial for identifying small targets like drones, especially in complex environments. While spatio-temporal models do exist, they are mainly attention-based (e.g. RT-DETR), thus they cannot run in real time on resource-constrained edge Neural Processing Units (NPU) such as NXP i.MX8M Plus.

This research internship addresses the challenge of building lightweight temporal feature fusion directly into 2D-convolutional backbones while maintaining real time latency (~30ms). The core theoretical challenge lies in learning to distinguish physical target motion from background motion induced by mobile/moving camera platforms, using only NPU-supported 2D operators. By adapting efficient video understanding techniques (such as Temporal Shift Modules or early feature fusion) to open-source 2D detectors (e.g., YOLOX), this project aims to capture multi-frame temporal context with near-zero overhead under strict hardware operator constraints (NXP eIQ, TFLite format).

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 (TFLite), ONNX
- **Frameworks & Libraries:** Python, PyTorch, OpenCV
- **Datasets & Data Pipelines:** USC Drone, Anti-UAV RGBT, RT-DETR (for automated hard-negative mining and benchmark curation)

3. Objectives

- **Literature & Operator Audit:** Review lightweight temporal modeling techniques (e.g., frame difference channels, TSM, feature-level shift, StreamYOLO) and ensure their operations are fully supported by the NXP NPU to guarantee deployment.
- **Dataset & Benchmark Pipeline Curation:** Assemble a multi-frame drone detection benchmark using open-source datasets (Anti-UAV RGBT, USC Drone). Build a specialized test set targeting single-frame false negatives using pre-annotations from large detection models (e.g. RT-DETR).
- **Architecture Design & Implementation:** Modify an open-source 2D detection backbone (e.g., YOLOX) to integrate zero-parameter or low-overhead temporal feature passing modules capable of learning temporal cues from consecutive frames.
- **Quantization & NPU Optimization:** Export the developed multi-frame architecture to TFLite, applying post-training quantization, potentially quantization-aware training and optimizing hardware execution on the NXP i.MX8M Plus NPU via the NXP eIQ pipeline.
- **Experimental Benchmarking:** Quantify trade-offs between frame-level mAP, temporal detection rate (P_d), false alarm rate per hour, and real-time hardware latency (FPS/ms) compared against single-frame

baselines under both fixed and mobile camera scenarios.

4. Where and when

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

5. References

- [1] J. Lin, C. Gan, and S. Han, "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, S. Divvala, R. Girshick, A. Fahradi, "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, S. Liu, Z. Li, X. Li, J. Sun, "StreamYOLO: Real-time Object Detection for Streaming Perception," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [4] Z. Ge, S. Liu, F. Wang, Z. Li, and J. Sun, "YOLOX: Exceeding YOLO Series in 2021," *arXiv preprint arXiv:2107.08430*, 2021,
- [5] Y. Quan, B. Kiefer, M. Messmer and A. Zell, "Lightweight Multi-Frame Integration for Robust YOLO Object Detection in Videos," *IEEE European Conference on Mobile Robots*, 2025,
- [6] Y. Zhao, W. Lv, S. Xu, J. Wei, G. Wang, Q. Dang, Y. Liu, and J. Chen, "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.

Multi-Temporal 3D Differential Change Detection for Volumetric Surface Monitoring

1. Context

Traditional two-dimensional (2D) change detection models compare planar aerial or satellite images frame by frame, thresholding pixel-wise discrepancies. However, in complex outdoor and industrial environments, 2D image differencing fails due to severe photometric and radiometric limitations, such as moving cast shadows, ambient lighting shifts, cloud cover, and seasonal vegetation changes, which generate massive false-positive alerts. Furthermore, perspective distortions (such as building lean and off-nadir tilt) and sub-pixel registration errors create edge-driven false positives along high-contrast boundaries. While elevation-based 2.5D Digital Surface Model (DSM) differencing reduces some perspective errors, it flattens vertical geometry and struggles on sloped or occluded structures.

This research internship addresses the challenge of building a 3D-native differential change detection framework that replaces photometric comparison with explicit geometric measurement. By operating directly on pre-aligned multi-temporal 3D representations, such as dense point clouds and 3D Gaussian Splatting (3DGS) scenes, the system decouples physical surface geometry (e.g., ground subsidence, material accumulation, structural warping) from transient visual appearance and illumination shifts. The core theoretical challenge lies in learning robust 3D-native distance metrics and directional consistency representations (e.g., 4D Gaussian differences or C-NeRF consistency fields) that isolate physical volumetric shifts without requiring manual distance thresholding.

2. Technologies

- **3D Differential & Distance Paradigms:** 4D Gaussian Splatting Change Detection, Gaussian Difference, C-NeRF (Directional Consistency NeRF), M3C2 (Multiscale Model-to-Model Cloud Comparison), Cloud-to-Cloud (C2C), Cloud-to-Mesh (C2M)
- **Alignment & Co-Registration:** Fast Point Feature Histograms (FPFH), Sample-Consensus Initial Alignment (SAC-IA), Generalized ICP (GICP), ENU Georeferencing
- **Frameworks & Libraries:** Python, PyTorch, Open3D, CloudCompare
- **Benchmarks & Datasets:** Urb3DCD, AHN-CD, Change3D, NeRFCD, KITTI / SemanticKITTI

3. Objectives

- **Literature & Algorithm Audit:** Review 3D-native point cloud differencing metrics (M3C2, C2C) and emerging 4D/3DGS radiance field change detection architectures (Gaussian-Difference, C-NeRF) to evaluate their trade-offs in noise robustness, computational latency, and threshold dependency.
- **Co-Registration & Alignment Pipeline:** Develop a robust multi-stage 3D co-registration workflow combining FPFH/SAC-IA coarse alignment with multi-stage GICP/ICP refinement and ENU georeferencing to guarantee centimeter-level spatial co-registration across temporal data captures.
- **3D Differential Change Engine:** Implement a 3D-native change detection worker operating on pre-aligned multi-temporal 3D Gaussians or point clouds to compute signed spatial displacement vectors (Z) and volumetric shifts while isolating true structural changes.
- **Photometric & Transient Filtering:** Integrate directional-consistency filters or learned feature classifiers (e.g., 4D Gaussian feature classification or SAM-based instance tracking) to filter out transient photometric variations, moving shadows, and seasonal vegetation noise.

- **Experimental Benchmarking:** Quantify performance (mIoU, Precision, Recall, F1-score, Chamfer distance) on open 3D change detection benchmarks (Urb3DCD, AHN-CD) measuring false-positive reduction compared against 2D image differencing baselines.

4. Where and when

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

5. References

- [1] U. Stilla and Y. Xu, "Change detection of urban objects using 3D point clouds: A review," *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 197, pp. 228–255, 2023.
<https://doi.org/10.1016/j.isprsjprs.2023.01.010>
- [2] R. Huang, H. Tao, B. Jiang, Q. Zhao, L. Wan, and Q. Guo, "C-NeRF: Representing Scene Changes as Directional Consistency Difference-Based NeRF," *IEEE Transactions on Image Processing*, vol. 34, pp. 7502–7513, 2025.
<https://doi.org/10.1109/tip.2025.3622374>
- [3] C. Pang, J. Wu, J. Ding, C. Song, and G.-S. Xia, "Detecting building changes with off-nadir aerial images," *Science China Information Sciences*, vol. 66, no. 4, 2023. <https://doi.org/10.1007/s11432-022-3691-4>
- [4] I. de Gélis, T. Corpetti, and S. Lefèvre, "Change Detection Needs Change Information: Improving Deep 3-D Point Cloud Change Detection," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 62, pp. 1–10, 2024.
<https://doi.org/10.1109/tgrs.2024.3359484>
- [5] B. Jiang, R. Huang, Q. Zhao, and Y. Zhang, "Gaussian Difference: Find Any Change Instance in 3D Scenes," *ICASSP 2025 - IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2025.
<https://doi.org/10.1109/icassp49660.2025.10890807>