EURANOVA

INTERNSHIP OFFERS 2025

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, software engineering, and data architecture. To do this, we offer counsel, R&D, and solutions.

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 and master thesis offers.

Explore Our Internship Opportunities

This document presents internship topics supervised by our consultancy and research & development 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, graph theory, artificial intelligence, data privacy, high-performance computing, etc.

We value continuous learning and teamwork, sharing our ideas and challenges, and we love to have a good time together. Our company culture is built on mutual enrichment, trust, care, respect, and knowledge sharing. 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.

Breaking Barriers: Tackling Bias in AI for Medical Image Segmentation

Context

Over the past few years, the integration of artificial intelligence (AI) systems for computer-aided diagnosis and image-based screening has become widespread within medical institutions. Consequently, the research community has dedicated significant efforts to developing accurate algorithms to support medical practitioners.

However, AI models trained using medical images for clinical tasks often exhibit bias in the form of disparities in performance between subgroups [1]. Biases can stem from gender, race, or age imbalances, as well as from acquisition protocols, scanner models, and annotation inconsistencies.

Biases in medical imaging datasets manifest in various forms, including prevalence bias, where certain attributes (such as specific diseases or demographic groups) are overrepresented, leading to skewed predictions. Presentation bias occurs when subgroups are either misrepresented or underrepresented, often due to disparities in healthcare access or diagnostic frequency. Annotation bias results from inconsistent labelling due to subjective human interpretation, where annotators' expertise or judgement may vary, influencing AI model predictions [9]. Despite the success of AI in medical imaging, little attention has been paid to the impact of such biases.

Recent advancements have specifically addressed this issue, highlighting how biased or unfair representation within datasets can significantly compromise the performance and fairness of AI systems [2]. These biases can severely affect the reliability of AI systems in critical tasks like cancer diagnosis or treatment planning [9]. Therefore, quantifying these biases and assessing their impact on model performance has become a critical milestone to ensure fairness and impartiality in network development.

This internship will focus on:

1. Bias Modelisation and Quantification: this step involves identifying and quantifying disparities generated via biases in medical image segmentation in a public medical dataset, the HECKTOR dataset, which focuses on PET/CT scans for head and neck cancer. The intern will explore bias sources such as scanner variations, gender, race, and disease prevalence.
2. Debiasing Algorithms: this step involves exploring and contributing to debiasing algorithms for medical image segmentation. The intern will improve upon existing debiasing techniques to improve model generalisation across diverse subgroups.

Objectives

The internship will focus on the following key objectives:

1. Preprocessing of the HECKTOR dataset: the intern will pre-process the HECKTOR dataset and prepare it for the deep learning segmentation model. To this end, they will use the KEDRO pipeline.
2. Cluster Identification and Bias Quantification: the intern will identify potential biases in the HECKTOR dataset generated due to underrepresented demographic groups, varying acquisition protocols... The intern will quantify biases using metrics such as class imbalance ratios, statistical tests, and visual tools (e.g., histograms, bias heatmaps) to identify disparities. To this end, the intern will gain expertise using visual tools such as Plotly, MLflow, …
3. Debiasing Methods: the intern will contribute to already developed and new debiasing techniques for medical image segmentation to ensure fair representation across all subgroups.

4. Fairness Evaluation: the intern will assess the impact of biases on machine learning algorithms and evaluate fairness metrics across different subgroups, providing insights into the effects of debiasing efforts.

Competencies

Throughout this internship, you will develop expertise in the following areas:

Knowledge

- Bias identification, quantification, and mitigation in AI for medical imaging
- Deep learning for debiasing medical image segmentation algorithms
- AI fairness methods and metrics
- Medical imaging modalities (PET, CT)

Technologies

- Python (PyTorch (Lightning), Kedro, MLflow, Pandas, Scikit-learn, Plotly, Kedro, MLflow)
- Jupyter Notebooks
- Docker
- Gitlab

Where & When

The internship will begin in February/March and last for six months, based in the R&D department in the Marseille office.

References

[1] Stanley, E. A. M. et al., “Towards objective and systematic evaluation of bias in medical imaging AI,” 2023.
[2] Ganz, M., Holm, S. H., Feragen, A., "Assessing Bias in Medical AI," 2023.
[3] Jones, C., Castro, D. C. S. Ribeiro, et al., "No Fair Lunch: A Causal Perspective on Dataset Bias in Medical Imaging," 2023.
[4] Piçarra, C., Glocker, B., "Analyzing Race and Sex Bias in Brain Age Prediction," 2023.
[5] Seyyed-Kalantari, L., et al., “Underdiagnosis bias of artificial intelligence algorithms applied to chest radiographs in under-served populations,” Nat. Med., 2021.

Synthetic Data in Computer Vision: Adaptation and Evaluation

Context

Data availability is one of the primary challenges when training computer vision models. Not only is there a general scarcity of data, but there are also specific scenes and situations for which data is entirely absent. For instance, labelled images of military vehicles are often difficult or impossible to obtain.

To address this issue, at Euranova, we have developed an internal toolset, hereafter referred to as "SITCOM", designed to generate synthetic datasets. This toolset blends classical 3D modelling techniques, which allow for precise control over the generated scenes, with generative AI-based refinement techniques, enhancing the realism and quality of the synthetic data.

While the data may be adequate to the human eye, this does not mean its performance in model training will be satisfactory. Indeed, we run into another common issue: domain gap. How to ensure the characteristics of the synthetic image are close enough to reality that a model trained on one will fare well on the other? How to quantify the gap that can exist between the two? How to assess the impact of performance on training? Our current approach uses a score-based method [2] where an external vision transformer model assesses the realism of the images, but this can be biased. Instead, we want to explore distribution-based approaches such as [1].

In order to further bridge this gap, we would also like to explore a new class of approach. At Euranova, we have drone pilots, and have the capability to scan existing environments to produce 3D models by photogrammetry. Gaussian Splatting methods constitute the new standard in terms of realistic rendering for pre-existing scenes. As such, we want to explore the use of Gaussian Splatting based on these scans to generate novel views of these environments for the purpose of generating synthetic data and, crucially, develop a pipeline to add arbitrary objects (cars, trees, etc.) to a scene scanned into a point cloud and render it with Gaussian Splatting.

As such, this internship will focus on three main key points:

- Implementation of certain missing functionalities in SITCOM.
- Evaluation of the domain and performance gap between synthetic and real data.
- Exploration of 3D Gaussian Splatting.

Objectives

This internship will have several objectives. They are presented below in their estimated chronological order.

1. In the existing pipeline, for generative AI, implement custom inpainting masks to keep certain objects and implement an API for the more recent models.
2. Strengthen the evaluation pipeline using new scores and study the impact of adding synthetic data to already implemented training strategies.
3. Develop a new feature: rendering novel views of existing scenes using Gaussian Splatting based on existing views acquired by drone flight.
4. Determine if we can easily combine Gaussian Splatting with existing mesh-based approaches, for example, to add a mesh into a point cloud (with the Gaussian Splat annotations).
5. Applying this to Euranova projects that will require synthetic data at the time.

Technologies

- PyTorch
- Diffusion models
- Gaussian Splatting
- 3D reconstruction (photogrammetry)
- 3D modelling (Blender)
- Computer vision foundation models (MobileNet, etc.)

Where and when

France (Marseille), from March to August 2025.

References

[1] Hyungtae Lee and Yan Zhang and Yi-Ting Shen and Heesung Kwon and Shuvra S. Bhattacharyya, "Exploring the Impact of Synthetic Data for Aerial-view Human Detection", (2024), https://arxiv.org/abs/2405.15203

[2] Max Ku, Dongfu Jiang, Cong Wei, Xiang Yue, Wenhu Chen - VIEScore: Towards Explainable Metrics for Conditional Image Synthesis Evaluation - ACL 2024 https://tiger-ai-lab.github.io/VIEScore/

Development And Deployment Of AI-Based Assistance Tools For The Automatic Analysis Of 3D High-Dimensional Medical Images

Context

In the recent era, deep learning has emerged as a pivotal tool for the automation of tedious tasks for high-dimensional medical image analysis, particularly in the sector of cancerous tumour segmentation [1] and tracking [2] in 3D scanner images. However, despite the undeniable success of deep learning-based technologies for medical image processing, two major limitations persist in bridging the gap between medical and technical expertise and enabling the adoption of these tools by physicians.

The first limitation is the need for large amounts of precise and exhaustive annotations of 3D images by medical experts for the development and evaluation of performant models. The annotation process is a time-consuming task, and despite the recent development of new annotation tools for AI applications [3], the tools available on the market do not always meet the technical requirements, especially for the tracking of tumours in medical images. Therefore, Euranova developed a custom annotation tool that was successfully used by physicians for the labelling of a first sample of data. The first objective of the internship will be to improve, optimise and automate the deployment of this annotation tool, based on the physicians' feedback and requirements.

The second limitation is the deployment of these tools in real-world environments. The complexity and diversity of the medical images require the development of robust and adaptable pipelines that can be deployed on a variety of infrastructures, particularly in different hospital environments. The second part of the internship will focus on developing a Minimum Viable Product (MVP) for the automatic segmentation and tracking of tumours in high-dimensional images that can be deployed in real-world environments. The development of this MVP will require connecting several pipelines for the processing of medical images, including already trained AI models, and developing an interface to visualise the different steps of the pipeline.

The intern will work closely with AI researchers and medical professionals, contributing to a project that combines cutting-edge technology with real-world medical challenges, ultimately enhancing the precision of tumour diagnosis and treatment. The trainee will have the opportunity to work with existing functional tools and also to develop end-to-end software.

Objectives

- Web Application Development and Deployment:
    - Develop the backend of the current Dash-based web application for the labelling of high-dimensional medical images to integrate new features based on the physicians' feedback and requirements. This task will, for instance, include the addition of a new 3D image modality.
    - Improve the UX by ensuring the application is user-friendly, ergonomic and efficient.
    - Improvement of the UI to give the tool a more modern appearance.
    - Automate and deploy the tool for rapid production and use as part of an annotation campaign.

- MVP for Tumour Segmentation and Tracking:
    - Create an end-to-end pipeline that connects several image processing pipelines and pre-trained AI models, ensuring seamless interaction and testing the robustness of the entire system.
    - Develop a modern and simple UI to visualise and highlight the underlying AI technology as well as the most important medical outputs (biomarkers).

    - Deploy the MVP using Docker and set up automated workflows with GitLab CI/CD for streamlined deployment and continuous updates.
    - Ensure that the software performs well on new data and that the pipelines can be adapted to different configurations or future improvements.

- Collaboration and Documentation:
    - Collaborate with AI researchers, medical professionals, and engineers to ensure the tools meet the required standards.
    - Document the development and deployment process, making it easily reproducible for future iterations.

Technologies

- Python: Kedro, Dash, PyTorch, Scikit-learn, SimpleITK, MLflow, Plotly
- Docker
- Git, GitLab, CI/CD

Where and when

France (Marseille), from February to September 2025.

References

[1] Hosny, Ahmed, et al. "Clinical validation of deep learning algorithms for radiotherapy targeting of non-small-cell lung cancer: an observational study." The Lancet Digital Health 4.9 (2022): e657-e666.

[2] Cai, Jinzheng, et al. "Deep lesion tracker: monitoring lesions in 4D longitudinal imaging studies." Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. 2021.

[3] https://monai.io/label.html

Sample-Efficient Intrinsically Explainable Reinforcement Learning Agent

Context

Amongst others, Reinforcement Learning (RL) suffers from two problems:

- Sample inefficiency: training takes forever to converge (the agent model has to "see" millions of times the same situations to robustly learn from them)
- Lack of interpretability: we have no idea what "motivates" an agent to make some moves at a given point in time, and on a more macro level, we don't understand its global behaviour.

The former is partially linked to the complexity of the (observed state, action) space. This problem has been solved in other ML fields by finding a semantic latent space where similar raw input points are projected onto the same point, simplifying downstream training from this embedding space. The explainability problem can also benefit from proper semantic space embedding and action space embedding, since it can allow a more hierarchical view of the agent actions and environment.

Objectives

In this context, the intern will have a few tasks to reach useful outcomes:

- Explore the state of the art of eXplainable Reinforcement Learning (XRL):
    - Review the different XRL techniques with a focus on global explainability;
    - Understand the limitations of current RL algorithms when dealing with explainability.
- Explore the start of the art of representation learning for RL:
    - Identify interesting approaches and test them on simple single-agent scenarios (e.g. OpenAI gym environments);
- Develop a new method that integrates embedding techniques into XRL approaches to improve sample efficiency and explainability;
    - Embedding techniques could be used both at the observation level and at the action level.
- Evaluate the new approach on a simple single-agent scenario. The evaluation should take into account training speed and stability, generalisation capabilities and explainability.

Technologies

- Single-agent Reinforcement Learning
- Deep learning (PyTorch)
- Representation learning
- Explainable Reinforcement Learning

Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2025.

References

[1] Fujimoto, Scott, et al. "For sale: State-action representation learning for deep reinforcement learning." *Advances in Neural Information Processing Systems* 36 (2024).

[2] Agarwal, Rishabh, et al. "Contrastive behavioral similarity embeddings for generalization in reinforcement learning." *arXiv preprint arXiv:2101.05265* (2021).

[3] Heuillet, Alexandre, Fabien Couthouis, and Natalia Díaz-Rodríguez. "Explainability in deep reinforcement learning." *Knowledge-Based Systems* 214 (2021): 106685.

[4] Lin, Zhengxian, Kim-Ho Lam, and Alan Fern. "Contrastive explanations for reinforcement learning via embedded self predictions." *arXiv preprint arXiv:2010.05180* (2020).

Light-Weight Multi-Object Tracking On Videos With Deep Learning Detection

Context

Multi-object tracking (MOT) is a critical challenge in computer vision, with broad applications in surveillance, autonomous driving, and robotics. In this use case, MOT aims to identify and track multiple objects in an image sequence over time, providing a continuous and robust identification of objects. Many approaches already exist; however, many of the best ones don't have a permissive licence for commercial applications.

This internship will focus on applying classical and state-of-the-art techniques for multi-object tracking in video. The intern will be responsible for developing algorithms for object tracking, based on Deep Learning object detection, as well as evaluating their performance and compute-efficiency on client and benchmark datasets.

Our specific goal is to develop robust tracking solutions for humans and vehicles in short video streams, with three key constraints:

1. (Near) Real-time tracking
2. Design for future deployment on embedded devices
3. Only use tools/data with permissive licences

Currently, our object detection and tracking rely on a licensed library. We aim to develop an in-house solution for these tasks. This internship will primarily focus on creating a tracking algorithm that meets the project requirements, with the potential to extend into object detection based on the intern's interests and project progress. While integrating the algorithms into embedded devices is beyond the scope of this internship, the focus will be on producing optimised algorithms for peak performance under edge device constraints.

Technologies

- Deep Learning (ONNX, TensorFlow Lite, PyTorch)
- Light-weight Object Detection
- Multi-Object Tracking: From Bayesian filter-based methods (e.g., Kalman, Particle), Joint Probabilistic Data Association Filter (JPDAF), to modern approaches (e.g., Tracktor, TrackFormer, ByteTrack, FairMOT)

Objectives

The primary goal of this internship is to develop probabilistic tracking approaches for multiple objects. Specific objectives include:

- Familiarising yourself with existing work on deep learning models for object detection and the commercial tracking algorithm;
- Developing and evaluating classical and state-of-the-art tracking algorithms for human and vehicle detection;
- Exploring solutions for handling challenging scenarios such as occlusions, clutter, and camera motion;
- Optimising tracking algorithms for (near) real-time performance on embedded devices;
- Contributing to a technical report that summarises the findings of the internship.

Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2025.

References

[1] G. Ciaparrone, F. L. Sánchez, S. Tabik, L. Troiano, R. Tagliaferri, F. Herrera, “Deep learning in video multi-object tracking: A survey”, 2019 – doi: 10.1016/j.neucom.2019.11.023

[2] P. Bergmann, T. Meinhardt, L. Leal-Taixe, “Tracking without bells and whistles”, 2019, doi: 10.1109/ICCV.2019.00103

[3] T. Meinhardt, A. Kirillov, L. Leal-Taixe, C. Feichtenhofer, “TrackFormer: Multi-Object Tracking with Transformers”, 2021 - doi: 10.48550/arXiv.2101.02702

[4] Y. Zhang, P. Sun, Y. Jiang et al., “ByteTrack: Multi-Object Tracking by Associating Every Detection Box”, 2022, doi: 10.48550/arXiv.2110.06864

[5] Y. Zhang, C. Wang, X. Wang et al., “FairMOT: On the Fairness of Detection and Re-Identification in Multiple Object Tracking”, 2021, doi: 10.48550/arXiv.2004.01888

[6] The Multiple Object Tracking Benchmark - https://motchallenge.net/

Evaluation of Graph RAG Strategies for Efficient Information Retrieval

Context

In recent months, Graph-based Retrieval-Augmented Generation (RAG) has emerged as a promising approach for improving the quality of information retrieval in AI systems. However, designing and implementing an efficient Graph RAG strategy involves numerous challenges and complexities. The goal of this internship is to explore and evaluate key strategies across two primary steps of the Graph RAG process: (1) graph construction and (2) information retrieval, with a focus on optimising performance and relevance.

The specific objectives are to assess the effectiveness of various techniques for constructing knowledge graphs and enhancing information retrieval before generating responses. The intern will investigate:

Graph Construction Techniques:

Several methods can be employed for extracting relationships and structuring a graph, including:

- Entity Relationship Extraction using NLP Tools: Utilising traditional NLP methods to identify and represent relationships between entities in text.
- Large Language Models (LLMs): Leveraging dedicated LLMs like RELIK [6] or frameworks like LlamaIndex [5] to construct knowledge graphs more efficiently through advanced AI techniques.

Information Retrieval Strategies:

Before generating a response, relevant context needs to be identified. There are multiple approaches for retrieving relevant information from a graph: [2,4]

- Graph Traversal for Context Extraction: Navigating the graph to extract relationships that can enhance the context for generation.
- Hybrid Graph Traversal and Vector Search: Combining graph traversal with re-ranking of information using vector search to optimise the quality of retrieved data.
- Vector Search on Graph Embeddings: Performing vector-based searches on graph embeddings and incorporating traversal techniques for improved retrieval of meaningful data. [7]

Objectives

1. Develop and define an evaluation framework for assessing different graph construction strategies.
2. Compare and contrast graph construction techniques, focusing on LLM-based methods versus traditional NLP approaches.
3. Design an evaluation strategy for measuring the effectiveness of information retrieval methods.
4. Experiment with graph traversal techniques to enhance context retrieval using graph embeddings, including the use of explainable graph embeddings to retrieve more meaningful and interpretable information.

By the end of the internship, the intern will have evaluated key strategies for graph construction and retrieval, and will propose improvements based on empirical findings. The intern will also have the opportunity to work with cutting-edge LLMs and advanced graph embedding techniques to solve real-world challenges in information retrieval and generation.

Technologies

- Graph DB, Kuzu
- Python
- Large Language Model
- Vector search
- Graph embedding

Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2025.

References

[1] Kuzu, the embedded graph DB https://github.com/kuzudb/kuzu

[2]GraphRAG with Kuzu https://github.com/kuzudb/graph-rag

[3] Xiaoxin He, Yijun Tian, Yifei Sun, Nitesh V. Chawla, Thomas Laurent, Yann LeCun, Xavier Bresson, Bryan Hooin. G-Retriever: Retrieval-Augmented Generation for Textual Graph Understanding and Question Answering. Arxiv https://arxiv.org/abs/2402.07630

[4] Graph RAG with Neo4Jhttps://neo4j.com/developer-blog/enhance-rag-knowledge-graph/

[5] https://docs.llamaindex.ai/en/stable/examples/query_engine/knowledge_graph_query_engine/

[6] Orlando, Riccardo, et al. "ReLiK: Retrieve and LinK, Fast and Accurate Entity Linking and Relation Extraction on an Academic Budget." Findings of the Association for Computational Linguistics ACL 2024. 2024.

[7] Gregory Scafarto, Madalina Ciortan, Simon Tihon, Quentin Ferre, Augment to Interpret: Unsupervised and Inherently Interpretable Graph Embeddings. In Proc. of The 15th Asian Conference on Machine Learning (ACML 2023), November 2023.

Consultant Assistant: A Centralised Pointer To Operational Knowledge

Context

Knowledge management is hard. Knowing what happened in the past, or at least who may know about what, is an everyday task in a consulting company. Each consultant sees a lot of different contexts, experiences many operational contexts, and solves many problems. Scale that to a ~200-employee company like Euranova, and knowledge becomes a gold mine to solve future technical problems.

Consultants onboarding on a new mission need to know from whom they can get help and pointers inside the company, and Business Managers regularly need to enrich commercial proposals with similar past experiences. In both cases, the problem is the same: given a rough description of a project (most likely in business terms and not with the proper technical names), get pointers to past experiences and resources to ramp up quickly on the associated technical challenges, as well as the names of people who have worked on similar tasks.

The data available takes multiple forms (free text info from consultants, slide decks presenting use cases, HTML data), which implies the need for both LLMs and VLMs. Some of those modalities are especially hard to integrate into a retrieval pipeline, which is going to spice things up for the intern, compared to classical retrieval pipelines, which can now be developed in a few lines of code.

In case there is enough time, an extension of the problem, which would have access to the internet, would also retrieve pointers about specific topics. In case representing the data as a knowledge graph is identified as necessary, some synergies will take place with the intern working on the GraphRAG-related topic.

Technologies

- Usage of LLMs (RAG systems, information representation)
- Usage of VLMs (QA based on slides)
- Online data scraping
- Python

Objectives

During this internship, the intern will:

1. Identify the different sources of data that should be included in the tool;
2. Survey the consultant population to gather their needs, identify non-functional constraints, and translate them into requirements;
3. Implement several retrieval systems, getting closer and closer to requirements and improving retrieval quality;
4. For each iteration of the system, evaluate its performance;
5. The system will be put in production and tested by the community.

Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2025.

References

[1] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, Douwe Kiela - Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks https://arxiv.org/abs/2005.11401

[2] Stephen E. Robertson et Karen Spärck Jones - Relevance weighting of search terms - Journal of the American Society for Information Science, vol. 27, no 3, mai-juin 1976, p. 129–146
https://dl.acm.org/doi/10.5555/106765.106783

[3] VLM-1 Cookbook https://github.com/autonomi-ai/vlm-cookbook

Novel View Generation From Inferred 3d Reconstruction On Images

Context

Machine learning needs some data to train with, and in a lot of industrial contexts, we are missing a proper amount of data to train models robustly. We may have the best architecture possible, but without data, we can't get reliable outcomes. In such settings, we can collect more data (which can be expensive) and label it (which can be expensive too).

Recent advances in synthetic data generation (in particular for computer vision tasks) offer new possibilities when it comes to producing data and labels. Approaches range from producing new samples from scratch (with some domain adaptation) using a relevant foundation model to re-rendering data points with a different perspective, with also the possibility to edit/alter samples in certain cases.

Data scarcity, in particular, is a problem when training image models based on drones' vision, or more generally, edge devices embarked on costly material. Euranova has developed a few drone-related assets throughout the years, and this issue comes back again and again.

There is always a sweet spot between producing synthetic samples with too much novelty (which may be far from the real sample distribution) and producing samples with too little of it (synthetic samples don't bring much to the discussion in such cases.). That is why in the case of drone vision, we are interested in the following pipeline, which we have the intuition could provide significant gains:

- From a drone sequence of photos, a 3D reconstruction is built using the latest advances on the matter;
- From this reconstruction, alternative drone trajectories are simulated and rendered.

By doing so in a proper way, segmentation labels on the original photo sequence will be propagated without additional labelling work to the newly rendered data points.

Technologies

- Deep learning on images (generation, segmentation)
- Gaussian splatting, Nerfs
- Python

Objectives

During this internship, the intern will:

- Study our current pipelines related to drone vision;
- Propose data augmentation pipeline(s) to generate new trajectories from existing data;
- Include domain adaptation strategies to correct the gap between real and synthetic data;
- Evaluate the impact of the generation strategy on the main task of interest.

Where and when

Belgium (Mont-Saint-Guibert), from March/April to August/September 2025.

References

[1] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis - 3D Gaussian Splatting for Real-Time Radiance Field Rendering https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/ - SIGGRAPH 2023

[2] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, Hengshuang Zhao - Depth Anything V2 https://arxiv.org/abs/2406.09414 - Neurips 2024

[3] Jonas Geiping, Micah Goldblum, Gowthami Somepalli, Ravid Shwartz-Ziv, Tom Goldstein, Andrew Gordon Wilson - How Much Data Are Augmentations Worth? An Investigation into Scaling Laws, Invariance, and Implicit Regularization https://arxiv.org/abs/2210.06441

[4] Hyungtae Lee and Yan Zhang and Yi-Ting Shen and Heesung Kwon and Shuvra S. Bhattacharyya, "Exploring the Impact of Synthetic Data for Aerial-view Human Detection", (2024), https://arxiv.org/abs/2405.15203