

EURV
NOVA

**INTERNSHIP
OFFERS
2024**

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.

Sustainable AI & Digital Services

1. Context

Digital technology has an environmental impact, increasingly contributing to the **carbon footprint and electricity consumption** each year [1].

Those in the ICT (Information and Communication Technology) sector play a pivotal role in shaping how technology affects society and the environment. They make the most of state-of-the-art technical knowledge to bring life to customers' great ideas with high-quality solutions. The **environmental impacts** of those solutions have to be quantified, limited and even anticipated. AI solutions that adhere to responsible AI practices take sustainability as an essential aspect of innovation success. [2,3]

In many cases, sustainability in data and AI coincides with the frugality of resources without compromising on quality. Indeed, computing resources are often a major expense in data projects, whether in terms of time or money. Integrating this kind of knowledge with usual technical skills into operational guidelines can help Euranova consultants and the broader AI developers' community create sustainable services.

Emerging legislation and regulations regarding the environmental impact of digital technology [4, 5] **are compelling the industry to make sustainability a core priority**. Moreover, the call for sustainability-aware projects is intensifying as it becomes a stipulated **requirement** in industrial R&D project tenders. It is increasingly common to request the development of strategies for sustainable development with minimal environmental impact as part of the solution.

Your goal during this internship is to **explore** literature and **craft** a set of essential knowledge for feeding Euranova's Responsible AI (RAI) service offer. The RAI framework includes ethics, privacy, transparency, security and other aspects that are crucial factors for the success of responsible AI projects. This internship aims to build on top of that by extending the framework with a sustainability pillar to help software and AI practitioners evaluate and mitigate the environmental impact of their digital activities.

2. Competencies

We expect you to learn and develop your skills in the following area:

- **Knowledge**
 - Data: data engineering & artificial intelligence
 - Ecodesign: between business and technical challenges
 - Cloud: architecture, cloud native & software engineering
 - Law & Regulation: environmental, digital
 - Knowledge in analytics and business process management
- **Technologies**
 - Frugal AI
 - Carbon Aware DevOps
 - Software & Web Development best practices

3. Objectives

Your mission is to deliver a report document summing up the essential set of processes a project has to go through, the minimal roles that need to be assigned, the basic training people will need to follow, and the list of state-of-the-art technical tools developers will need to use for designing sustainable AI and digital services. To obtain the desired deliverable, you may proceed from the following roadmap:

1. Build a comprehensive state of the art about sustainable computing; the legislations and company obligations, the risks, the technologies, the processes and the advantages.
2. Sum up the core principles for sustainable computing by keeping in mind the framework has to be flexible and actionable. Extract the practical essence of every information source.
3. Look into existing tools and knowledge in and out of Euranova.
4. Come up with selection criteria for technical tools (for measuring, reducing/limiting the impact) and business roles and processes (for anticipating the impact).
5. Propose a minimal technical toolbox and guidelines for operationalizing the framework.
6. Elaborate a framework relating to ICTs core businesses (data governance, data engineering, data science, software engineering, cloud native, architecture), establishing for each how the impact can be measured and what can be done to reduce it.
7. Evaluate the commercial value of such a sustainable computing framework.

4. Where and when

The internship will begin in **February/April 2024** and will last for **5 to 6 months**, in the R&D Department in the **Marseille** office. It will be supervised by experts from France and Belgium who may provide advice and feedback on both technical and business aspects.

5. References

- [1] ADEME & Arcep. (2022) "[Évaluation de l'impact environnemental du numérique en France et analyse prospective.](#)" Observatoire des impacts du numérique.
- [2] Luccioni, Alexandra Sasha, and Alex Hernandez-Garcia. (February 2023) "[Counting carbon: A survey of factors influencing the emissions of machine learning.](#)" *arXiv preprint arXiv:2302.08476*.
- [3] Renee Cho. (June 2023) "[AI's Growing Carbon Footprint.](#)" State of the Planet. Columbia Climate School.
- [4] Lois et Décrets. (2021). "[LOI n° 2021-1485 du 15 novembre 2021 visant à réduire l'empreinte environnementale du numérique en France \(1\).](#)" Journal officiel de la République française.
- [5] European Parliament. (June 2023). "[Amendments on the proposal for a regulation of the European Parliament and of the Council on laying down harmonised rules on artificial intelligence.](#)" Artificial Intelligence Act

Siamese Networks For Tumour Tracking In Medical Imaging

1. Context

In the recent era, deep learning has emerged as a pivotal tool for the diagnosis and follow-up of various medical conditions, particularly in the sector of cancerous tumour detection and tracking [1, 2, 3]. Despite the undeniable success of deep learning in the field of medical image processing and diagnosis, one major limitation still persists: the need for large amounts of annotated training data that enables their high performance. The objective of this project is to leverage cutting-edge deep learning-based technologies in order to address the challenge of tumour tracking between PET/CT baseline and follow-up scans. Tumour tracking is particularly challenging due to changes in the acquisition process between baseline and follow-up scans, as well as the patient's position and the unpredictable spread of tumours in the body.

The objectives of this project are twofold:

1. The candidate must explore new deep learning-driven solutions that enable fully supervised deep learning methods for accurate matching and tracking of tumours between paired baseline PET/CT scans and follow-up scans [4].
2. The second objective is to investigate semi- and weakly supervised solutions that can mitigate the need for labour-intensive tumour-annotated datasets [5, 6].

The expected outcomes are carried out thanks to the availability of an in-house private dataset, which was obtained through a collaboration with a public entity. In addition to the private dataset, we also have access to a public dataset that is similar in nature to the private one. These datasets provide a robust foundation for training, testing, and validating our proposed deep-learning models.

2. Competencies

We expect you to learn and develop your skills in the following area:

- **Knowledge:**
 - Deep Learning for Object Tracking
 - Semi-, weakly, and fully supervised learning
 - Medical imaging (PET, CT)
- **Technologies:**
 - Python (PyTorch (Lightning), Scikit-learn, Plotly, Kedro, Mlflow)
 - Jupyter Notebooks
 - Conda, Docker
 - Gitlab.

3. Objectives

The internship will have the following objectives:

- Establish a baseline fully supervised deep learning model (Siamese Network) capable of accurately tracking tumours given a limited-sized fully annotated dataset.
- Train and evaluate previously developed deep-learning models on a new dataset.
- Explore state-of-the-art approaches for weakly and semi-supervised object tracking.

- Generate weakly supervised labels and train the proposed framework.
- Conduct a comprehensive evaluation of the developed methods relative to performance, efficiency, and scalability, as well as the ability of weakly supervised methods to reduce the need for full annotations.

The results of this internship could be presented to the community as a talk, a report or a scientific paper. To conduct this work, the intern will participate in discussions with physicians.

4. Where and when

The internship will begin in March/April 2024 and will last 6 months, in the R&D department in the **Marseille** office.

5. References

- [1] M. Ganeshkumar, V. Sowmya, E.A. Gopalakrishnan, K.P. Soman, Chapter 10 - Unsupervised deep learning-based disease diagnosis using medical images, Editor(s): Akash Kumar Bhoi, Victor Hugo C. de Albuquerque, Parvathaneni Naga Srinivasu, Gonçalo Marques, In Intelligent Data-Centric Systems, Cognitive and Soft Computing Techniques for the Analysis of Healthcare Data, Academic Press, 2022, Pages 203-220, ISBN 9780323857512, <https://doi.org/10.1016/B978-0-323-85751-2.00011-6>.
- [2] Li-Yin Ye, Xiao-Yan Miao, Wan-Song Cai, Wan-Jiang Xu, Medical image diagnosis of prostate tumor based on PSP-Net+VGG16 deep learning network, Computer Methods and Programs in Biomedicine, Volume 221, 2022, 106770, ISSN 0169-2607, <https://doi.org/10.1016/j.cmpb.2022.106770>.
- [3] Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. "U-net: Convolutional networks for biomedical image segmentation." International Conference on Medical image computing and computer-assisted intervention. Springer, Cham, 2015.
- [4] Grama D, Dahele M, van Rooij W, Slotman B, Gupta DK, Verbakel WFAR. Deep learning-based markerless lung tumor tracking in stereotactic radiotherapy using Siamese networks. Med Phys. 2023 May 23. doi: 10.1002/mp.16470. Epub ahead of print. PMID: 37219823.
- [5] Mengyan Zhang, Cong Wang, Li Cai, Jiyun Zhao, Ye Xu, Jiacheng Xing, Jianghong Sun, Yan Zhang, Developing a weakly supervised deep learning framework for breast cancer diagnosis with HR status based on mammography images, Computational and Structural Biotechnology Journal, Volume 22, 2023, Pages 17-26, ISSN 2001-0370, <https://doi.org/10.1016/j.csbj.2023.08.012>.
- [6] Jinxi Xiang, Xiyue Wang, Xinran Wang, Jun Zhang, Sen Yang, Wei Yang, Xiao Han, Yueping Liu, Automatic diagnosis and grading of Prostate Cancer with weakly supervised learning on whole slide images, Computers in Biology and Medicine, Volume 152, 2023, 106340, ISSN 0010-4825, <https://doi.org/10.1016/j.compbiomed.2022.106340>.
- [7] Gatidis S, Kuestner T. A whole-body FDG-PET/CT dataset with manually annotated tumor lesions (FDG-PET-CT-Lesions) [Dataset]. The Cancer Imaging Archive, 2022. DOI: 10.7937/gkr0-xv29

Perturbation-based XAI Method Modular Framework Enhancement

1. Context

Explainable AI (XAI) is a pivotal element in the development of AI systems, particularly given the recent impetus from regulations such as the GDPR or the AI Act. XAI allows data scientists to gauge the pertinence of their models and provides end-users with insightful revelations, enabling them to interact knowledgeably with model predictions. Among the various XAI methodologies, those based on perturbation have gained widespread adoption. These methods adhere to a principle of altering elements within the original input and observing the subsequent changes in the model's output, thus identifying the important features.

The past few years have seen the introduction of numerous explicability methods (RISE, LIME, SHAP, MP, etc.). Often, these are compiled into libraries for convenience; however, they tend to be presented merely as compilations, lacking in modularity. Notably, scant attention has been paid to dissecting their shared characteristics and structures. At Euranova, we have put forth a theoretical decomposition into which most perturbation-based explicability methods can fit. This decomposition involves breaking down methods into constituent steps (Exploration, Perturbation, Evaluation, Aggregation), plus a Memory system, allowing for a mix-and-match approach.

This work has led to the creation of a library named Muppet, which facilitates the implementation of these methods following the aforementioned decomposition. Due to its modular nature, the fundamental components of each method can be reused in the construction of other methods, or even in the creation of new innovative methods with relative ease.

As of this writing, the preliminary phase of this endeavour has concluded, and an initial paper has been submitted. Our current ambition is to augment the library to encompass the full spectrum of popular perturbation-based interpretability methods and to embed benchmarking functions that can swiftly evaluate the performance of newly developed variants within Muppet.

The initial segment of the internship will heavily involve Python development, as the primary objective is to expand an existing library. Nonetheless, this expansion will be consistently anchored to a theoretical and mathematical objective and will require a scientific analysis of the methods in order to re-implement them properly, not merely a code transposition. A precise list of features to be incorporated in order to complete the library has been prepared.

Should these goals be attained early enough, the intern would then be tasked with devising a novel interpretability method leveraging the Muppet framework. The intern's scientific perspective in order to create such a new method would be both solicited and valued. If necessary, existing preliminary work that has been carried out internally on several potential new methods could be leveraged in order to guide the reflection. The creation of this new XAI method would serve to validate the design of Muppet, ensuring that the development and evaluation of a fresh XAI method is both easy and straightforward within the Muppet framework.

2. Technologies

- Pytorch
- Python
- GitLab

3. Objectives

- Implement a Benchmark of methods using the Quantus library (faithfulness, etc.).

- Implement various methods: OptiCAM, genetic algorithms, HiPE, LIME, tree structures, etc. within the Muppet modular framework. This implementation is not a simple code transposition but involves analysing and understanding existing explicability methods to map them into the theoretical framework.
- Implement multimodality using the existing Premise component (i.e. tabular data and not just images).
- Documentation and code improvements.
- Write a scientific paper detailing the new implementations and the potential new methods and method variants developed using the tool, with an evaluation made by the benchmarking tools implemented.

4. Where and when

The internship will begin in March/April 2024 and will last 6 months, in the R&D department in the **Marseille** office.

5. References

[1] Lundberg, S.M., Lee, S.I.: *A unified approach to interpreting model predictions*. Advances in neural information processing systems 30 (2017)

[2] Arrieta, A.B., Díaz-Rodríguez, N., Del Ser, J., Bennetot, A., Tabik, S., Barbado, A., García, S., Gil-López, S., Molina, D., Benjamins, R., et al.: *Explainable artificial intelligence (xai): Concepts, taxonomies, opportunities and challenges toward responsible ai*. Information fusion 58, 82–115 (2020)

[3] Hedstrom, Anna and Weber, Leander and Krakowczyk, Daniel and Bareeva, Dilyara and Motzkus, Franz and Samek, Wojciech and Lapuschkin, Sebastian and Hohne, Marina M-C, *Quantus: An explainable ai toolkit for responsible evaluation of neural network explanations and beyond*, Journal of Machine Learning Research 2023

[4] Kokhlikyan, Narine and Miglani, Vivek and Martin, Miguel and Wang, Edward and Alsallakh, Bilal and Reynolds, Jonathan and Melnikov, Alexander and Kliushkina, Natalia and Araya, Carlos and Yan, Siqi and others, *Captum: A unified and generic model interpretability library for pytorch*, arXiv preprint arXiv:2009.07896, 2020

[5] Zhang, Hanwei and Torres, Felipe and Sicre, Ronan and Avrithis, Yannis and Ayache, Stephan, *Opti-CAM: Optimising saliency maps for interpretability*, arXiv preprint arXiv:2301.07002, 2023

Early Detection of Concept Drift

1. Context

The essential conceit of Machine Learning is that past data can be used to make predictions about yet-unseen data. But what happens when this assumption is not the case? Distribution shifts are a potential way for this problem to arise and are surprisingly common in ML. They can come from a change in the data distribution. But they may also result from a difference in the decision rules, for instance, if an external disruption (such as the COVID-19 pandemic in our case) changed the interactions at play.

A Machine Learning algorithm that does not account for these changes may have poor accuracy when operationally deployed. As such, for critical applications, it is a fundamental problem in terms of safety. Euranova has already encountered this COVID-19 effect in temporal data several times, and being able to detect it, characterise it and correct it is of major importance. It is usually denoted as a concept shift, where the relation between features and targets in historical data changes over time. This means the same input variables result in different outcomes in each situation.

Formally, considering the inputs to a model (X) and its output (Y), we have $P(X, Y) = P(Y|X)P(X) = P(X|Y)P(Y)$. A distribution shift happens when any of these component distributions change over time. To account for distribution shifts, it is necessary to detect and quantify them precisely and seek a way to reverse or even neutralise them beforehand, possibly through projection. This necessitates modelling the data's generative process and, as such, detecting shifts between conditional distributions $P(Y|X)$.

In this internship, we are interested in exploring the literature on concept drift detection in time series data. Once a robust methodology for comparing data distributions has been found, we will study the problem of detecting concept shifts in Air Traffic Management data recorded before, during and after the COVID-19 outbreak. In addition, we want to explore the moment of triggering the alarm for concept drift in online settings where new measurements are gathered over time [1].

2. Technologies

- Knowledge in Python and ML libraries: PyTorch, numpy
- Git/Gitlab

3. Objectives

- Review of state-of-the-art methods on Concept drift detection.
- Listing and familiarising with existing Python libraries for concept drift detection.
- Apply the algorithms to a motivating use case on the Eurocontrol air traffic data.
- Develop a library (or improve an open-source one) with the implementation of the distribution shift characterization
- Propose a new adaptive method for early detection of concept drift.
- Write a scientific paper to summarise your findings.

4. When and Where

The internship will begin in February 2024 and will last 6 months, in the R&D department in the **Marseille** office.

5. References

- [1] Alexis Bondu, Youssef Achenchabe, Albert Bifet, Fabrice Clerot, Antoine Cornuejols, Joao Gama, Georges Hebrail, Vincent Lemaire, and Pierre-Francois Marteau. 2022. Open challenges for Machine Learning based Early Decision-Making research. SIGKDD Explor. Newsl. 24, 2 (December 2022), 12–31.
- [2] Ditzler, G., Roveri, M., Alippi, C., Polikar, R., & Rosen, G. (2015). Learning in nonstationary environments: A survey. IEEE Computational Intelligence Magazine, 10(4), 12-25.
- [3] Bifet, A., & Gavalda, R. (2009). Learning from time-changing data with adaptive windowing. In Proceedings of the 2009 SIAM international conference on data mining (pp. 443-454).
- [4] Schmidt, P., Wollmer, M., Krajewski, J., & Eyben, F. (2018). Detecting concept drift in affective computing systems: A feature and data distribution-based approach. In International Conference on Affective Computing and Intelligent Interaction (pp. 273-279).
- [5] Gama, J., Zliobaite, I., Bifet, A., Pechenizkiy, M., & Bouchachia, A. (2014). A survey on concept drift adaptation. ACM Computing Surveys (CSUR), 46(4), 44.

Language Powerhouse: Exploring LLM Potential

1. Context

Large Language Models (LLMs) have emerged as a transformative force in natural language processing. These models, such as GPT-3, LLAMA, MISTRAL and others, have been trained on large amounts of text data, enabling them to accurately generate human-like text. They have impacted many different fields. For example, LLMs can power chatbots in customer service, improve healthcare outcomes by helping to manage and access healthcare information, or even provide productivity help in general by streamlining document analysis and review in a variety of contexts.

Developing or customising a Large Language Model tailored to a project's need is a challenging endeavour. It can incur major costs for development, computational resources, and data acquisition. It further necessitates the acquisition of specific expertise. Ethical considerations, such as the introduction of biases, are also potential problems. All of this necessitates meticulous strategic planning. As such, we seek to investigate these issues, and explore the possibility of tailoring and applying an LLM to a business project.

In this internship, we will explore the process and constraints of developing a specialised LLM for a given project, from beginning to end. The LLM may potentially be used in any aspect of the project which is suitable. With the help, enthusiasm and technical knowledge of the intern, we seek to acquire a deeper understanding of the capabilities and limitations of LLMs, as well as develop and test a methodology for their integration into a concrete project.

2. Activities

This internship is a dedicated exploration of harnessing the potential of Large Language Models (LLMs). It encompasses a comprehensive study of LLM applications and entails a hands-on involvement in tasks throughout the entire LLM lifecycle. This includes data collection, cleaning, normalisation, tokenisation, thoughtful model selection, fine-tuning for task-specific optimisation, rigorous evaluation, seamless deployment, and the ongoing commitment to continuous monitoring and maintenance. By participating in these multifaceted responsibilities, the intern will gain an in-depth understanding of LLMs and their practical applications.

3. Objectives

- Explore the entire LLM lifecycle through one project: review the state-of-the-art, cover the implementation and evaluation, and deploy;
- Define a comprehensive set of metrics to gauge a model's readiness for real-world applications;
- Deploy LLM models across a variety of applications.

4. When and where

The internship will begin in March/April 2024 and will last 6 months, in the R&D department in the **Marseille** office.

5. References

- [1] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, Illia Polosukhin, Attention is All You Need (2017) NIPS
- [2] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan and al. (2020) Language Models are Few-Shot Learners, NeurIPS
- [3] Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, Gideon Mann, BloombergGPT: A Large Language Model for Finance (2023)
- [4] Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, LoRA: Low-Rank Adaptation of Large Language Models, 2021
- [5] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas and al. Training Compute-Optimal Large Language Models, 2022

Advanced State Estimation with Neural Network

1. Context

In the context of autonomous navigation, vehicle location is essential. This location can be obtained via GNSS when it is available and accurate. However, in the absence of GNSS - due to GNSS jamming, areas with poor GNSS coverage, indoor environments, etc. - we rely on other methods to successfully perform the localisation task. To compensate for the absence of GNSS, mobile location is based on other modalities: computer vision, IMU, and aerial imagery. In turn, these modalities are not always available, and their measurements are noisy. In practice, some of the states in a system cannot be measured directly, and the measurements are not accurate enough due to sensor uncertainty. Traditionally, location estimation (referred to as state estimation in the following) is based on filters, notably the Kalman filter. The application of a simple Kalman filter has already been explored, but its performance in borderline cases is unsatisfactory. In particular, when the measurement error is neither unimodal nor Gaussian.

2. Technologies

- Docker
- Python
- Git / GitLab

3. Objectives

The aim of this internship is to apply data-driven state estimation for information fusion in the context of autonomous navigation. The trainee will have to familiarise himself with the notions of classical state estimation and propose a data-driven architecture to solve the problems posed by Kalman filters. Two applications are possible within the framework of the course: ego-centric estimation and tracking estimation. The choice of application will be discussed with the candidate.

The first use case is a mission to transport goods in an area without GNSS coverage using an autonomous ground vehicle. The ground robot must be able to locate itself using its cameras, its IMU and a satellite image of the region. The navigation zone is unknown to the robot. In this case, the IMU, camera and satellite data must be merged to help localise it. The estimation problem is complex because of the ambiguous nature of visual recognition. Visual localisation sometimes offers up to 4 possible locations. Conventional unimodal filters cannot handle this situation. The objective will be to propose a data-based estimation of position and orientation by taking advantage of the power of neural networks.

The second use case is for aerial tracking of a land vehicle. A drone must work with the driver of a land vehicle. The drone follows the vehicle and returns an aerial view to the driver. In several cases, the transmission between the vehicle and the drone is cut off, and the drone no longer receives the vehicle's position. We have, therefore, set up a system for predicting the position of the ground vehicle based on historical positions and road maps. Our system uses a Kalman filter and a hidden Markov model. The current solution does not manage intersections and does not consider the driver's behaviour to predict his trajectory. The aim will be to propose a data-based trajectory prediction system using the power of neural networks.

4. Where and when

The internship will begin in March/April 2024 and will last 6 months, in the R&D department in the **Marseille** office.

5. References

[1] R. E. Kalman, "A New Approach to Linear Filtering and Prediction Problems," *Journal of Basic Engineering* 82, no. 1 (March 1, 1960): 35–45, <https://doi.org/10.1115/1.3662552>.

[2] Mohinder Grewal and Angus Andrews, "Kalman Filtering: Theory and Applications," January 1, 1985.

[3] Norbert Wiener, *Extrapolation, Interpolation, and Smoothing of Stationary Time Series* (The MIT Press, 1964).

[4] Ahmed S. Zamzam, Xiao Fu, and Nicholas D. Sidiropoulos, "Data-Driven Learning-Based Optimization for Distribution System, State Estimation," *IEEE Transactions on Power Systems* 34, no. 6 (November 2019): 4796–4805, <https://doi.org/10.1109/TPWRS.2019.2909150>.

Synthetic Data Editing

1. Context

Synthetic data editing has emerged as a leading solution for obtaining GDPR-compliant data like "MidJourney", "DALL-E", and "Stable Diffusion", which are currently relatively popular tools in the field [1][2][3][4][5][6].

By producing highly realistic data, synthetic datasets maintain closeness to the original data's distribution, making them valuable for various machine learning tasks. One of the primary attractions of synthetic data is the balance it strikes between data utility and privacy [7] – ensuring data is both useful and private, especially in the absence of actual data.

This internship aims to refine synthetic data editing techniques to enhance the diversity in datasets, particularly regarding ethnicity and gender. This will involve transforming data attributes, such as morphing male characteristics into female ones or modifying physical traits indicative of one ethnicity to reflect another. Improving the diversity of datasets holds immense importance, especially when such datasets are employed to train models used in critical real-life scenarios. This guarantees the effectiveness of these tools across a broad spectrum of people.

Please note this internship is conducted in collaboration with a company working in the defence industry.

2. Competencies

We expect you to learn and develop your skills in the following area:

Knowledge:

- Image processing/Computer vision (at least notions).

Technologies:

- Python.
- Git / GitLab.
- ML libraries (PyTorch).
- Data science tools (Jupyter Notebook).

3. Objectives

In order to respond to the initial problem of finding a solution to improve diversity in a dataset without having to manually re-collect data, a non-exhaustive list was proposed:

- Make a rigorous study of the state of the art for synthetic data editing approaches.
- Search for a dataset and annotate it if data on gender, ethnicity, etc. are missing, to have sufficient material to train and evaluate future model(s).
- Implement a pipeline to train each model correctly.
- Implement a test pipeline to evaluate each model correctly.
- Obtaining model(s) with a good performance to obtain equality synthetic data.
- Implement a GUI to facilitate the utilisation of the model.

4. Where and when

The internship will begin in January 2024 and will last 6 months, in the R&D department in the **Mont-Saint-Guibert** office (Belgium).

5. References

- [1] Fan, Liyue. "A Survey of Differentially Private Generative Adversarial Networks." (2020).
- [2] Yang Li, Michael Purcel, Thierry Rakotoarivelo, David Smith, Thilina Ranbaduge, Kee Siong NG; "Private Graph Data Release: A Survey", 2022
- [3] Dena Magdy. "The Use of Artificial Intelligence Art Generator MidJourney in Artistic and Advertising Creativity." 2023.
- [4] Aditya Ramesh and Prafulla Dhariwal and Alex Nichol and Casey Chu and Mark Chen. "Hierarchical Text-Conditional Image Generation with CLIP Latents", 2022
- [5] Lee, Seongmin & Hoover, Benjamin & Strobel, Hendrik & Wang, Zijie & Peng, ShengYun & Wright, Austin & Li, Kevin & Park, Haekyu & Yang, Haoyang & Chau, Polo. "Diffusion Explainer: Visual Explanation for Text-to-image Stable Diffusion", 2023.
- [6] Jun-Yan Zhu and Taesung Park and Phillip Isola and Alexei A. Efros. "Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks", 2020.
- [7] Zhao, Benjamin Zi Hao, Mohamed Ali Kâafar and Nicolas Kourtellis. "Not one but many Tradeoffs: Privacy Vs. Utility in Differentially Private Machine Learning." Proceedings of the 2020 ACM SIGSAC Conference on Cloud Computing Security Workshop (2020): n. pag.

Reinforcement Learning and Large Language Models: Synergies and Opportunities

1. Context

Reinforcement Learning (RL) and Large Language Models (LLMs) are two distinct yet active fields of research in Artificial Intelligence (AI), and both have witnessed important achievements in recent years.

RL is a branch of AI where one or more agent(s) learn(s) to make sequences of decisions by interacting with an environment to maximise some notion of cumulative reward. RL can be used in various applications, including smart energy management or robotics [1]. It even reached superhuman performance levels in several strategy games, such as Go [2] and StarCraft II [3]. Despite these impressive results, many challenges remain to be solved in order to deploy versatile and robust RL solutions in the industry.

LLMs are advanced models capable of understanding and generating human language. These models, exemplified by GPT-3.5 [4] or LLaMA [5], have revolutionised natural language processing tasks, enabling applications in translation, content creation, and even dialogue systems such as ChatGPT. LLMs learn patterns and structures within languages, allowing them to generate coherent and contextually relevant text or even code. Yet, several issues related to the training of these models and the hallucination phenomenon are still under investigation.

Recently, researchers have started to investigate the potential synergies between these fields. For instance, a major challenge in RL is the design of the reward function. Indeed, unexpected and nonsensical results may quickly arise from a badly designed reward function. In this perspective, recent work tested the idea of using an LLM as a proxy to translate a desired behaviour formulated in natural language into the code of a reward function well-suited for the RL task [6, 7].

Another problem to tackle in RL is the exploration/exploitation tradeoff [1]. The state space of an RL task is often huge. Think of all the move combinations in a Chess game, for example. Therefore, RL agents must learn when to try new moves to discover possibly better strategies and when to exploit the current strategy to maximise the reward. Here, LLMs can aid RL agents in exploring the state space intelligently. By generating diverse and informative queries, LLMs could guide the agent's exploration strategy, leading to faster and more effective learning in novel or complex tasks.

On the other hand, RL is also beneficial to LLMs. A popular example is that of ChatGPT, which leverages Reinforcement Learning from Human Feedback (RLHF) for training and fine-tuning [8]. Generally, RL might be applied to improve the performance of LLMs and enable them to generate more accurate, coherent, and contextually appropriate responses [9].

Overall, this internship offers a unique opportunity to delve into an active line of research at the intersection of RL and LLMs, investigating innovative ways to leverage their combined power for smart sequential decision-making and natural language understanding.

2. Technologies

- Reinforcement learning
- Large Language Models
- Python programming
- Git / Gitlab

3. Objectives

The goals of this internship include, but are not limited to:

- Reviewing the literature of RL and LLMs to identify potential synergies.
- Extending existing frameworks. As a starting point, the intern can leverage the recent Eureka framework to design reward functions for other RL environments, such as MPE [10] or Star Craft II [11]. More information on Eureka and some demonstrations: <https://github.com/eureka-research/Eureka>
- Formulating novel ideas and hypotheses on how one technology might benefit the other.
- Implementing, testing and validating these ideas in a rigorous pipeline.
- Identifying the limitations of the approach and recommending areas of improvement to future researchers.

4. Where and when

The internship will begin in January 2024 (flexible) and will last 6 months, in the R&D department in the Mont-Saint-Guibert office (Belgium).

5. References

- [1] Sutton, Richard S., and Andrew G. Barto. *Reinforcement learning: An introduction*. MIT press, 2018.
- [2] Silver, David, et al. "Mastering the game of Go with deep neural networks and tree search." *nature* 529.7587 (2016): 484-489.
- [3] Vinyals, Oriol, et al. "Grandmaster level in StarCraft II using multi-agent reinforcement learning." *Nature* 575.7782 (2019): 350-354.
- [4] Brown, Tom, et al. "Language models are few-shot learners." *Advances in neural information processing systems* 33 (2020): 1877-1901.
- [5] Touvron, Hugo, et al. "Llama: Open and efficient foundation language models." *arXiv preprint arXiv:2302.13971* (2023).
- [6] Ma, Yecheng Jason, et al. "Eureka: Human-Level Reward Design via Coding Large Language Models." *arXiv preprint arXiv:2310.12931* (2023).
- [7] <https://www.marktechpost.com/2023/07/20/researchers-from-stanford-and-deepmind-come-up-with-the-idea-of-using-large-language-models-llms-as-a-proxy-reward-function/> (Consulted on October 31st 2023, 11am CET)
- [8] Stiennon, Nisan, et al. "Learning to summarize with human feedback." *Advances in Neural Information Processing Systems* 33 (2020): 3008-3021.
- [9] Gulcehre, Caglar, et al. "Reinforced self-training (rest) for language modeling." *arXiv preprint arXiv:2308.08998* (2023).

[10] Mordatch, Igor, and Pieter Abbeel. "Emergence of grounded compositional language in multi-agent populations." *Proceedings of the AAAI conference on artificial intelligence*. Vol. 32. No. 1. 2018.

[11] Samvelyan, Mikayel, et al. "The starcraft multi-agent challenge." *arXiv preprint arXiv:1902.04043* (2019).

AI & Data Risk Management

1. Context

In the contemporary landscape of rapidly advancing technology, the integration of Artificial Intelligence (AI) and its synergy with data is transformative, driving innovation and efficiency. However, this integration also introduces challenges and risks that organisations must navigate adeptly. AI, powered by machine learning algorithms, is heavily reliant on diverse datasets for training and decision-making. Managing the associated risks involves addressing issues such as algorithmic bias and data privacy concerns, and ensuring the security of sensitive information.

The convergence of AI and data mandates a comprehensive risk management approach that encompasses ethical considerations, regulatory compliance, and the development of transparent and explainable AI models. Euranova, therefore, proposes a Responsible AI solution that operationalises AI ethics and guides the trustworthy design, development, deployment, and use of AI systems in practice while evaluating and monitoring for machine learning risks.

This internship aims to build on the existing responsible AI and risk assessment approach to dive deeper into AI and data risk management frameworks, tools, methodologies, and best practices. The goal is to develop a risk management framework to identify, assess, manage, and mitigate risks associated with the utilisation of data and AI within an organisation. The primary purpose of such a framework is to systematically address the challenges and uncertainties inherent in data-driven and AI-enabled processes and applications.

2. Competencies

We expect you to learn and develop your skills in the following areas:

- Knowledge of analytics, artificial intelligence, and business
- Knowledge of data management, data & AI governance, privacy, and risk management
- Knowledge of business process management

3. Objectives

Your mission is to develop a framework for data and AI risk management. The internship will undergo the following phases:

1. **Research & Analysis:** We recommend dividing this process into two steps:
 - a. Compile an elaborate review of *AI risk assessment* frameworks, tools, methodologies, and best practices.
 - b. Extend the research into data & AI risk management approaches.

The output of both steps is a state-of-the-art report that could provide inputs for an academic publication/survey.

2. **Modelling:** develop a risk management framework by
 - a. Integrating Responsible AI concepts and the AI risk assessment approaches.
 - b. Designing risk mitigation strategies and controls to address issues such as data privacy, algorithmic bias, and security vulnerabilities.

- c. Establishing mechanisms for ongoing monitoring and assessment of risks as the organisation's data and AI landscape evolves.
- d. Linking to data governance policies, procedures, and roles for managing data and AI models throughout their lifecycle.

The output of the modelling phase is a technical report.

4. Where and when

The internship will begin in **February/ March 2024** and will last for **4 to 6 months**, in the R&D department in **Tunis**.

Responsible AI Standardisation & Compliance

1. Context

In an evolving regulatory, ethical, and business landscape, organisations aiming to derive value from AI are faced with the challenge of adhering to ethical principles, gearing up for regulatory compliance, and aligning business processes and technical practices with international standards.

To respond to these challenges, Euranova proposes its own Responsible AI framework which provides insights and guidance on how best to adapt to the changing AI landscape and ensure that organisations have the correct policies, procedures, and technical capabilities in place to guarantee ethical and responsible outcomes from AI adoption. These outcomes are ensured through the development of transparent, robust, privacy-preserving, and risk-aware machine learning models by organisations that are accountable and compliant.

The goal behind our work on responsible AI is to assist organisations in the trustworthy development, deployment, and use of AI.

On the other hand, various international organisations are working on developing international standards for responsible AI. For example, the International Organization for Standardization (ISO) has established working groups to create standards related to AI ethics, bias mitigation, and other aspects of responsible AI.

In Europe, the EU Commission is putting forth the European AI Act, a legislative proposal to regulate the development and use of AI as part of the EU's broader strategy to establish a trustworthy and human-centric approach to AI.

This internship aims to extend the existing RAI framework with a focus on exploring the RAI standardisation and regulatory landscape and benchmarking these efforts to our framework with the purpose of extending it based on the findings.

2. Competencies

We expect you to learn and develop your skills in the following areas:

- Knowledge of analytics, artificial intelligence, and business
- Knowledge of data management, data & AI governance, privacy, and risk management
- Knowledge of business process management

3. Objectives

This project is organised around upgrading the responsible AI framework.

1. **Research & Analysis:** At this phase, the objective is to:
 - a. Conduct structured and thoroughly documented research into AI and RAI standardisation frameworks.
 - b. Review the requirements of the European AI Act and other AI legal frameworks.
 - c. Explore ethical RAI tradeoffs and links to machine learning tradeoffs for responsible AI.
2. **Benchmarking:** Use the findings of step 1 to map the reviewed standards and regulatory requirements to Euranova's responsible AI framework.
3. **Evaluation:** Evaluate the existing tools and resources to validate the conformity with the AI Act.

4. **Modelling:** Upgrade the RAI framework.

4. Where and when

The internship will begin in **February/ March 2024** and will last for **4 to 6 months**, in the R&D department in **Marseille or Tunis**.

Privacy Assessment of Graph-Generative Models

1. Context

In the last year, generative AI has impressed greatly with its sheer power and applicability in a tremendous number of use cases, ranging from chatbots to image generation, copywriting, and many more. Unfortunately, it has also created worries in the sector, especially with claims of hallucinations, copyright infringement, biased outputs, etc.

Regulatory frameworks are being put in place, especially in the EU, to protect and safeguard people's rights, especially privacy, against these emerging technologies. In this internship, we offer the opportunity to contribute to a long-ongoing project about privacy evaluation in generative models applied to graph data.

Usually, privacy scores are made in one of two ways: either you simulate an attack, and your privacy score is the attack's success rate, or you train differentially private models that blur the models in a way that protects the privacy of the training samples.

However, both these techniques are inherently flawed; the former is limited to known attacks, which doesn't protect against new types of attacks: in practice, it's virtually impossible to assess all known attacks. The latter was shown to introduce much more noise than was actually required to guarantee privacy, thus drastically reducing the model's performance.

This internship aims to finish a project that is well underway and showing promising preliminary results, but that needs finishing touches to be published. There are also opportunities to extend the work to use cases that are closer to the constraints actually facing industries today.

2. Competencies

We expect you to know and/or be willing to familiarise yourself with the following knowledge:

- Pytorch
- Graph Machine Learning
- Deep Learning
- Privacy evaluation

3. Objectives

The objectives of the project are structured as follows:

1. Gain missing knowledge about graph machine learning and previous work at Euranova on privacy for graph generative models.
2. Finish the existing paper to make it ready for submission at a conference (and submit it!)
3. Extend the privacy framework to one of the following domains:
 - a. Temporal Graphs
 - b. The case where the generative model only has a single (large) input graph.

4. Where and when

The internship will begin in January 2024 (flexible) and will last 6 months, in the R&D department in the **Mont-Saint-Guibert** office (Belgium).

LLM for Data & AI regulation

1. Context

The European Union's data regulatory framework is becoming increasingly intricate, posing significant challenges in compliance and understanding. The landscape includes various regulations such as the Data Act, e-Privacy, Data Governance Act, General Data Protection Regulation (GDPR), and other specific data exchange regulations. This complexity often overwhelms even the most experienced legal teams, particularly considering the vast array of jurisprudence texts. Consequently, there is a pressing market need for tools that can offer comprehensive insights into these legal constraints, tailored to specific use cases.

This internship focuses on leveraging the capabilities of Large Language Models (LLMs) to address this challenge. LLMs have demonstrated remarkable efficiency in processing and interpreting extensive text data, making them ideal for this application. The project aims to develop a sophisticated tool using a ready-to-productize LLM, either open-source or sourced from a partner vendor. This tool will integrate legal texts and jurisprudence to facilitate the understanding of the legal landscape for specific use cases.

2. Competencies

We expect you to know and/or be willing to familiarise yourself with the following knowledge:

- Python
- Machine Learning
- Deep Learning
- LLM basic understanding

3. Activities

This internship is a dedicated exploration of harnessing the potential of Large Language Models (LLMs) in this legal context. It encompasses a comprehensive study of LLM legal applications and entails a hands-on involvement in tasks throughout the entire LLM lifecycle. This includes data collection, cleaning, normalisation, tokenisation of legal content (provided to the internship student). Euranova will provide a set of "legal use cases" with specific questions for their particular context. The student may have to fine-tune task-specific optimisation but also provide rigorous evaluation and a seamless deployment. By participating in these multifaceted responsibilities, the intern will gain an in-depth understanding of LLMs and their practical applications.

4. Objectives

- a. Integration of Legal Texts: The intern will work on integrating various EU legal texts and jurisprudence into the LLM. This involves not just data ingestion but also ensuring that the model understands the context and relevance of different legal documents.
- b. User Query Processing: The tool will be designed to process queries from legal professionals, delivering clear, concise explanations of the legal context relevant to specific use cases. This feature aims to demystify the complex web of EU data regulations for its users.

- c. Impact Analysis: A critical functionality of this tool will be to analyse how specific jurisprudence impacts one or more regulations. The intern will develop methods for the LLM to identify and explain these relationships and their potential effects on various scenarios.
- d. Usability and Accessibility: The intern will also focus on making the tool user-friendly for legal professionals, ensuring easy access and interpretation of complex legal information.

From a scientific perspective:

1. Explore the entire LLM lifecycle through one project: review the state-of-the-art, cover the implementation and evaluation, and deploy;
2. Constitute the legal documentation repository
3. Study state-of-the-art methods for integrating this content within the LLM knowledge base
4. Setup evaluation method
5. Proceed to a rigorous evaluation of the LLM
6. Study state-of-the-art method for crossing LLM extraction (jurisprudence and legal text) to identify impacts

5. Where and when

The internship will begin in January 2024 (flexible) and will last 6 months, in the R&D department in the **Mont-Saint-Guibert** office (Belgium).