Internship opportunity in multi-modal learning for forest monitoring

Deadline to apply: June 7th 2025, 11.59pm Anywhere on Earth.
Questions? arthur.ouaknine@mila.quebec with [intern!] in the subject line
Submit your application here*.

Duration:
Six-month, full-time internship at Mila in research focusing on the intersection of forest monitoring and machine learning. In collaboration with David Rolnick and Étienne Laliberté.

Project description:
Applications of machine learning and computer vision to high-resolution drone data are revolutionizing forest monitoring [1]. Individual tree detection, segmentation, and classification are essential tasks in domains such as forest conservation and carbon stock estimation [2, 3]. Orthomosaics produced from photogrammetry provide exceptional details for deep learning detection and segmentation algorithms [4, 5], while recent research demonstrates that incorporating digital surface models (DSMs) alongside RGB imagery substantially enhances both traditional methods and foundation models like SAM [6]. Building on these advances, this project will investigate whether raw photogrammetric point clouds—initially used to produce DSMs—can improve supervised learning representations for individual tree detection and segmentation using recent open-access datasets [7, 8]. The research will be then pursued in benchmarking these methodologies for individual tree height estimation using solely RGB-derived information (orthomosaics and 3D photogrammetric point clouds).

Role Responsibilities:

- Collaborate on an open source project[1] to pre-process georeferenced orthomosaics and dense photogrammetric point clouds from photogrammetry.
- Setting up a georeferenced pre-processing pipeline to downsample and discretize dense 3D photogrammetric point clouds at multiple resolutions and scales with large volumes (from one to several tens of millions of points) using Rasterio and Laspy.
- Design, construct, and validate experiments using PyTorch and sublibraries (e.g. PyTorch Geometric).
- Structure a strategy of experiments to evaluate methods including ablation studies as well as performances on downstream tasks.
- Define evaluation metrics and structure routines to evaluate remote sensing based methods.
- Benchmarking existing deep learning methods for individual tree crown detection, segmentation and classification (convolutional and transformer-based architectures) from drone orthomosaics.

[1] https://github.com/hugobaudchon/geodataset

- Develop and implement point cloud-derived feature extraction techniques and multi-modal fusion strategies to enhance supervised deep learning frameworks for individual tree crown detection, segmentation, and classification.
- Extend existing approaches for individual tree height estimation in a multi-modal and multi-task framework.
- Propose data analysis and visualization of large volumetric orthomosaics and 3D point clouds.
- Collaborate with other team members to leverage existing code, documentation and datasets.
- Help writing documentation and scientific content, such as academic journals or conference papers.

Logistics:

- Duration: 6 months, paid full-time.
- Flexible start date, preferably early June 2025.
- Location: In-person at Mila in Montreal, Canada, with the option to start remotely.

Candidate Profile:

- Availability to work full-time for a minimum of 4 months, preferably 6 months, between June and December 2025.
- A Master or PhD level of expertise in machine learning is required (not the diploma).
- An intermediate level of expertise in computer vision and PyTorch is required (see questions below).
- Experience in at least one of the following areas is required:
    - Instance segmentation with remote sensing data (preferably) or natural images.
    - Multi-modal fusion and multi-task learning with remote sensing data (preferably) or natural images.
    - 3D point cloud modeling using deep learning.
- Experience with remote sensing data (from satellites and / or drones), geometric deep learning, forest monitoring or plant biology is a plus.
- Publication(s) in computer vision (CVPR, ICCV, ECCV) or machine learning (NeurIPS, ICLR, ICML) conferences or journals in related fields (TPAMI, JMLR) is a plus, but not required.

Please go to the form* to submit your cv and short responses:

- A link to **your cv online, or upload pdf**.
- Provide your response to **any 2 of the following 3** questions (total maximum 1 page):
    a. Can you describe a project you've completed using PyTorch? What was the challenge, and how did your solution address it?
    b. How have you used transfer learning in PyTorch to leverage pre-trained models for a specific task? Can you discuss the process and the outcomes?

c. Have you contributed to any open-source PyTorch projects or developed custom PyTorch modules? Please describe your contribution and learning from the experience.

*When submitting your materials online, it won't work from an incognito mode browser.

Thanks! We will contact prospective interviewees within two weeks of the application deadline. If you do not hear back, thank you for your interest... not all applicants will be contacted.

References:

[1] A. Ouaknine, T. Kattenborn, E. Laliberté, and D. Rolnick. OpenForest: a data catalog for machine learning in forest monitoring. Environmental Data Science, 4:e15, 2025.

[2] N. I. Bountos, A. Ouaknine, I. Papoutsis, and D. Rolnick. FoMo: Multi-Modal, Multi-Scale and Multi-Task Remote Sensing Foundation Models for Forest Monitoring. Proceedings of the AAAI Conference on Artificial Intelligence, 39(27):27858–27868, Apr. 2025.

[3] G. Reiersen, D. Dao, B. Lütjens, K. Klemmer, K. Amara, A. Steinegger, C. Zhang, and X. Zhu. ReforesTree: A Dataset for Estimating Tropical Forest Carbon Stock with Deep Learning and Aerial Imagery. Proceedings of the AAAI Conference on Artificial Intelligence, 36(11):12119–12125, June 2022.

[4] B. G. Weinstein, S. Marconi, M. Aubry-Kientz, G. Vincent, H. Senyondo, and E. P. White. DeepForest: A Python package for RGB deep learning tree crown delineation. Methods in Ecology and Evolution, 11(12):1743–1751, Dec. 2020.

[5] J. G. C. Ball, S. H. M. Hickman, T. D. Jackson, X. J. Koay, J. Hirst, W. Jay, M. Archer, M. Aubry-Kientz, G. Vincent, and D. A. Coomes. Accurate delineation of individual tree crowns in tropical forests from aerial RGB imagery using Mask R-CNN. Remote Sensing in Ecology and Conservation, 9(5):641–655, Oct. 2023.

[6] M. Teng, A. Ouaknine, E. Laliberté, Y. Bengio, D. Rolnick, and H. Larochelle. Assessing SAM for Tree Crown Instance Segmentation from Drone Imagery. ICLR 2025 ML4RS workshop, Mar. 2025.

[7] M. Cloutier, M. Germain, and E. Laliberté. Influence of temperate forest autumn leaf phenology on segmentation of tree species from UAV imagery using deep learning. Remote Sensing of Environment, 311:114283, Sept. 2024.

[8] I. Lefebvre and E. Laliberté. UAV LiDAR, UAV Imagery, Tree Segmentations and Ground Measurements for Estimating Tree Biomass in Canadian (Quebec) Plantations. FRDR, July 2024.