

Leonardo Mullins

Machine Learning Engineer

 (512) 555-0184

 leonardo.mullins@gmail.com

 linkedin.com/example/leonardomullins

 Austin, TX 78701

STRENGTHS

-  **Production ML Judgment**
Chose practical reconstruction methods after comparing SfM, MVS, SLAM, and COLMAP outputs with field needs.
-  **Debugging Mindset**
Traced a stubborn ingestion failure through Python services and logging; peers soon sought help on similar issues.
-  **Clear Communication**
Presented model failures in plain language, helping product and field partners make confident release decisions.
-  **Research Curiosity**
Tested emerging vision methods with disciplined experiments, then kept useful findings in shared technical notes.
-  **Collaborative Delivery**
Git reviews and thoughtful handoffs made platform integration smoother; teammates valued reliable production follow-through.

SKILLS

Python

PyTorch

TensorFlow

OpenCV

NumPy

Computer Vision

Object Detection

Semantic Segmentation

Instance Segmentation

3D Reconstruction

SfM

MVS

SLAM

COLMAP

Git

LANGUAGES

English

Native

Spanish

Intermediate

SUMMARY

Machine Learning Engineer with five years of hands-on experience building **computer vision** systems for visual inspection, geospatial imaging, and construction technology. Develops Python, PyTorch, TensorFlow, OpenCV, and NumPy solutions for **object detection**, **semantic segmentation**, **instance segmentation**, and 3D reconstruction. Turns aerial, spherical, and 360 imagery into dependable signals for field operations, schedule analysis, and progress monitoring. Brings practical experience debugging large-scale pipelines, evaluating emerging methods, improving inference reliability, and moving prototypes into production. Works comfortably with research, product, platform, and field teams, communicating findings in plain language. Ready to bring curiosity, sound engineering judgment, and customer-focused ML delivery to CanyonSight Technologies.

EXPERIENCE

Machine Learning Engineer

BlueRidge Vision Labs  January 2025 - Present  Austin, TX

Owns computer vision development for a construction progress monitoring product, guiding model research, data workflows, production integration, and technical communication across platform, product, and field operations teams.

- Built **PyTorch** detection models for construction assets, improving review quality across difficult site imagery.
- Trained **instance segmentation** models for materials and structures, clarifying errors caused by occlusion and lighting.
- Reworked **OpenCV** and **NumPy** preprocessing, standardizing calibration, tiling, and metadata checks before inference.
- Prototyped **COLMAP** SfM and MVS workflows, supporting reliable site-level **3D reconstruction** views.
- Compared reconstruction quality with SLAM approaches, helping engineers select practical production methods.
- Packaged inference services with platform engineers, adding Git review checks for dependable handoffs.
- Presented failure analyses to product managers and field specialists, guiding release decisions.

Computer Vision Engineer

CivicFrame Analytics  January 2023 - December 2024  Chicago, IL

Developed production-oriented TensorFlow and PyTorch workflows that converted jobsite imagery into labeled observations for schedule and progress analysis, partnering with product, infrastructure, and construction specialists.

- Developed **TensorFlow** workflows that converted jobsite imagery into observations for schedule analysis.
- Trained **semantic segmentation** models for concrete, framing, and mechanical work areas.
- Investigated ingestion and batching failures through **Python** services, improving diagnostic logging.
- Prepared large datasets with NumPy and OpenCV, shortening weekly experiment turnaround.
- Evaluated detection methods through precision-recall analysis and class-level error reviews.
- Reviewed reconstruction outputs with construction specialists, connecting model quality with field decisions.
- Contributed maintainable modules through **Git** branches, peer reviews, documentation, and release retrospectives.

Machine Learning Research Associate

MapForge Research  September 2021 - December 2022  Madison, WI

Conducted supervised applied research in aerial computer vision, building reproducible experiments, curated datasets, baseline models, and introductory 3D workflows for review by graduate researchers and industry partners.

MY CAREER



- Machine Learning Engineer at BlueRidge Vision Labs (1.7 Years)
- Computer Vision Engineer at CivicFrame Analytics (1.9 Years)
- Machine Learning Research Associate at MapForge Research (1.2 Years)

- Conducted **object detection** research for **aerial imagery** using reproducible Python experiments.
- Created image segmentation baselines, supporting class-level comparisons across held-out validation images.
- Quality-checked multi-class datasets, resolving annotation inconsistencies before model evaluation.
- Implemented CNN, PyTorch, and TensorFlow baselines for transfer-learning comparisons.
- Tested OpenCV and NumPy normalization, cropping, and augmentation across image viewpoints.
- Built COLMAP workflows for **SfM** and **MVS**, documenting reconstruction artifacts and camera limitations.
- Presented **research** findings at a showcase, earning faculty recognition for clear **technical communication**.

PROJECTS

Construction Reality Mapping 📅 2025

Created a research prototype that combined 360 imagery, object detection, and COLMAP reconstruction for site progress review. Clear visual outputs helped nontechnical reviewers connect model behavior with field questions.

Aerial Segmentation Benchmark 📅 2022

Built a reproducible benchmark for aerial image segmentation with PyTorch, TensorFlow, OpenCV, and NumPy. Structured error reviews made model comparisons easier for research collaborators.

LEADERSHIP & AWARDS

- 2025 Engineering Impact Award, BlueRidge Vision Labs, for production computer vision delivery.
- 2022 Research Showcase Recognition, MapForge Research, for clear presentation of aerial ML findings.
- 2024 Peer Leadership Recognition, CivicFrame Analytics, for dependable code reviews and team support.

EDUCATION

Master's Degree in Computer Science

University of Wisconsin–Madison 🎓 GPA: 3.8 📅 2021 📍 Madison, WI

Coursework: Machine Learning, Computer Vision, Deep Learning, 3D Computer Vision, Data Mining

CERTIFICATIONS

- TensorFlow Developer Certificate 📅 2024
- AWS Certified Machine Learning – Specialty 📅 2025

TECHNICAL SKILLS

- **Programming Languages:** Python, SQL, Bash
- **Deep Learning Frameworks:** PyTorch, TensorFlow, Keras
- **Computer Vision Libraries:** OpenCV, NumPy, SciPy
- **Detection and Segmentation:** Object Detection, Semantic Segmentation, Instance Segmentation
- **3D Vision:** SfM, MVS, SLAM
- **3D Reconstruction Tools:** COLMAP, Camera Calibration, Point Clouds
- **Image Data:** Aerial Imagery, Spherical Imagery, 360 Imagery
- **Machine Learning Methods:** CNNs, Transfer Learning, Data Augmentation
- **Data Engineering:** Large-Scale Data Processing, Image Tiling, Metadata Validation
- **Model Evaluation:** Precision-Recall, Intersection over Union, Error Analysis
- **Software Development:** Git, Code Review, Technical Documentation
- **Cloud Machine Learning:** AWS, Model Deployment, Inference Services

PROFESSIONAL AFFILIATIONS

- Member, Association for Computing Machinery, Special Interest Group on Artificial Intelligence.

- Member, IEEE Computer Society, participating in computer vision and machine learning discussions.

ADDITIONAL INFORMATION

Work Status : Authorized to work in United States. No sponsorship required.

REFERENCES

AVAILABLE ON REQUEST