CV & ML Research Engineer at
INSEER - Remote

- Developed and fine-tuned detection and segmentation models, 2D pose estimation, 2D-to-3D reconstruction, vision transformers, and preprocessing motion-capture data for training and evaluation.

- Produced POCs, implemented post-processing refinement pipelines, and solved edge cases in bounding-box tracking and pose estimation, including occlusion handling, person re-identification, ID switches, pose-jitter correction in occluded frames, and developed privacy-aware face and person blurring algorithms.

- Improved reliability of ergonomic metrics by refining skeleton alignment, joint-angle, and torque computation algorithms.

- Sparsified and benchmarked multiple YOLO, collaborated with Keymakr Labs to generate human-annotated bounding boxes, resulting in a 15% mAP@0.5 improvement and ~2x in inference throughput.

- Collaborated with the CEO to create an interactive Excel Export feature of the complete platform.

- Optimized ML pipeline GPU and CPU utilization and reduced storage redundancy on AWS S3, reducing inference latency by 60-70% and compute costs by 50%.