

Job Title: AI Engineer — Computer Vision, Edge AI & Healthcare Compliance

2D/3D AI Models · SLM/RAG Edge Deployment · QA & Regulatory Affairs — 3DGait & Beyond

Company: CareCam <https://carecam.ai>

Location: Singapore (Hybrid), **Type:** Full-time

Experience level: 1 to 3 years. This role is open to early-career engineers.

Why This Role Is Different

CareCam is a medtech company building 3DGait, an AI-powered Software as a Medical Device (SaMD) already used in clinical trials to help clinicians assess and manage patients with mobility challenges. The AI you build here doesn't sit in a research notebook, it runs in real clinics, on real patients, and has to hold up to both clinical scrutiny and regulatory review.

You'll work directly with our founding team including our CTO and senior engineers on the next phase of Carecam – our new product Vis-Frailty and improving 3DGait: richer 2D/3D AI models, and bringing small language model (SLM) and retrieval-augmented generation (RAG) capabilities to the edge, closer to the point of care.

Project Background

3DGait has been successfully launched and is now used in clinical trials to support patient screening and rehabilitation monitoring. As adoption grows, we're pushing more of our AI stack out of the cloud and to the edge: running 2D/3D vision models and SLM+RAG pipelines directly in the clinic to cut latency, reduce dependence on connectivity, and keep sensitive patient data closer to where it's generated. That shift raises the bar: models need to be accurate, small enough to run on constrained hardware, and backed by the QA and regulatory documentation a healthcare product demands.

What Problems Are We Trying to Solve?

- **Model accuracy:** Continuously improving 2D/3D computer vision model performance so clinical decision support is something a clinician can actually trust.
- **Edge-constrained deployment:** Making SLM + RAG pipelines small, fast, and reliable enough to run at the clinic edge, not just in a well-resourced cloud environment.
- **Regulatory-grade evidence:** Producing the validation, testing, and documentation trail QA and regulatory affairs need to support SaMD claims (e.g. FDA, HSA, ISO 13485/62304-aligned processes).
- **Data privacy at the edge:** Keeping patient data secure and compliant (HIPAA/PDPA) when inference and retrieval happen outside a centrally-controlled cloud environment.
- **Consistency across environments:** Making sure models behave the same way across hospitals, devices, and patient loads.

What You'll Actually Work On

- **Build the models:** Design, train, and iterate on 2D/3D AI models (computer vision / gait & movement analysis) that power 3DGait's core clinical capability.
- **Deploy to the edge:** Build and deploy SLM-based RAG pipelines optimized to run at the edge - model compression, quantization, and latency/memory tuning for constrained clinical hardware.

- **Own model QA:** Design test protocols, benchmarking suites, and accuracy/performance metrics that hold up to clinical scrutiny, not just a notebook metric.
- **Partner on regulatory affairs:** Contribute technical documentation, validation evidence, and risk assessments aligned with SaMD requirements (FDA/HSA, ISO 13485/62304) - working alongside our clinical and regulatory leads, not carrying RA solo.
- **Keep the pipeline healthy:** Maintaining the CI/CD and cloud deployment pipeline (AWS/GCP) that ships your models.
- **Work close to the domain:** Collaborate closely with backend, mobile, and clinical teams so what you build integrates cleanly into the product clinicians actually use.
- **Document as you go:** Write clear notes on what you built, tested, and validated.

What Makes a Great Fit

- **Must have:** Strong Python, with hands-on experience building and training 2D/3D computer vision or deep learning models (e.g. PyTorch/TensorFlow).
- **Must have:** Familiarity with retrieval-augmented generation (RAG) pipelines. Coursework, side projects and self-directed learning all count.
- **Nice to have:** Experience deploying models to edge/embedded hardware (quantization, ONNX, CoreML, TensorRT, or similar) rather than cloud-only inference.
- **Nice to have:** Prior exposure to a regulated industry (healthcare/medtech) — QA test design, or contributing to regulatory documentation for a SaMD or similar product.
- **Nice to have:** Familiarity with Docker, basic CI/CD, and a cloud platform (AWS/GCP) — supporting skills, not the core of the role.
- **Mindset:** Genuine curiosity about AI/ML and health tech, comfortable working close to regulatory rigor rather than treating it as someone else's problem, and a proactive, ship-it attitude.

What you do not need: You do not need prior medtech or regulatory experience, and you do not need to have shipped a model to edge hardware before. If you have the computer vision fundamentals and the appetite to learn, we will teach you the rest here.

What We Offer

- **Real clinical impact:** Your models run in a live SaMD used with actual patients, not a research shelf.
- **Direct founder access:** Regular, hands-on time with Carecam's Tech lead and senior engineers.
- **A rare cross-section:** Cutting-edge AI (2D/3D vision, SLM/RAG at the edge) + medtech regulatory rigor.
- **Full ownership:** Small team, high trust; your model design and QA decisions ship.
- **A steep, supported learning curve:** You will be stretched deliberately, with the CTO and senior engineers close enough to unblock you when you need it. This is a role to grow in, not to arrive fully formed into.
- **Flexible, hybrid working arrangements:** We care about output, not seat time.

How to Apply

Send your resume, a brief cover letter on why this intersection of AI, edge deployment, and healthcare compliance interests you, and (if you have one) a link to relevant model or project work, to enquiry@carecam.ai.

Subject line: "AI Engineer Application – [Your Name]".

Join us at CareCam and help build the future of healthcare technology.