

About Steg.AI

Steg.AI develops AI-powered watermarking technology to protect and authenticate digital media. Our invisible watermarks are designed to be imperceptible to humans while remaining robustly detectable by our proprietary models.

Founded in 2019, Steg.AI is an NSF- and investor-backed startup based in Orange County, California, with a team of 5 PhDs in computer vision.

We're hiring a **Machine Learning Engineer** to help push the frontier of AI for watermarking, steganography, and media provenance.

Job Description

You will take state-of-the-art deep learning models for steganography and own their journey to production — optimizing and deploying them on desktop and cloud platforms. As part of the ML research team, you will serve as the critical link between cutting-edge research and customer-facing products, working closely with the engineering team.

Scope of Responsibility

- Own the productionization of novel steganography models
- Design and maintain inference pipelines for watermarking workloads
- Collaborate with the engineering team to integrate models into cloud platforms
- Benchmark inference performance and track regressions
- Translate research findings into robust, maintainable production code

Required Qualifications

- Proficiency with PyTorch and experience exporting models via ONNX
- Hands-on experience deploying deep learning models to cloud platforms
- Proficiency in Python and C/C++
- Strong collaboration and communication skills to work across research and engineering teams

Bonus Qualifications

- Familiarity with steganography, watermarking, or media provenance technologies
- Experience with image and video compression algorithms and codecs
- Familiarity with MLOps tooling: model registries, CI/CD for ML, and monitoring
- Computer Vision or ML publications or contributions to open-source frameworks

To apply, send your resume/cv to careers@steg.ai