

Senior Software Engineer

AI-Powered Construction Intelligence

Location: Washington, DC Metro Area (DMV) – **Preferred**

Experience: 5-8 years

Type: Full-time, Hybrid (1 day/week in-person)

Role Overview

AnchorGrid is building AI-powered estimating tools that are transforming how construction contractors work. We're eliminating the tedious manual processes that currently consume 40-60% of estimators' time, allowing them to focus on strategic decisions rather than data entry.

As a Senior Software Engineer, you will lead the algorithm design and model training pipeline for one or more construction “divisions” (industry categories, such as HVAC or plumbing) to extract structured data from complex PDF plans and specifications. You'll be part of a small, focused team shipping features weekly and seeing immediate impact on real users' workflows.

What You'll Build

Your work will be **80% Python backend development** focused on:

- **Document intelligence systems** that parse construction drawings and specifications using computer vision and OCR
- **AI-powered data extraction pipelines** that transform messy PDFs into structured, actionable data
- **Celery workers** for async document processing and job orchestration
- **FastAPI services** with WebSocket support for real-time updates
- **Computer vision algorithms** for fixture detection, room identification, and spatial analysis using text extractors and OpenCV

The remaining 20% involves frontend integration work (TypeScript/React) to connect your backend services to user-facing verification interfaces.

Primary Responsibilities

- Develop and maintain document processing pipelines using text extractors, OpenCV, and AI vision models

- Build FastAPI endpoints and Celery workers for async document analysis
- Implement computer vision algorithms for object detection and spatial relationship analysis
- Integrate with LLM APIs (OpenAI, Anthropic, Mistral) for intelligent text extraction and classification
- Design and optimize PostgreSQL schemas for construction domain data
- Work directly with beta customers to understand pain points and validate solutions
- Write clean, tested code following domain-driven design patterns
- Participate in weekly in-person sessions for architecture planning and pair programming

Required Qualifications

- **Based in DC Metro Area (DMV)** – able to commute for weekly in-person meetings - Preferred.
- **5-8 years of professional experience** with Python in production environments
- **Tech lead experience** - experience leading a significant technical project subsystem
- **Strong understanding of OCR and document processing** – you know the difference between Tesseract, cloud OCR APIs, and when to use each
- **Infrastructure as code** – comfortable writing Terraform configurations, deploying containerized services, and working with Docker in production
- **Experience with AI/ML in real applications** – you've integrated OpenAI/Claude APIs, worked with vision models, or built ML-powered features
- **Computer vision fundamentals** – understanding of image processing, object detection, and coordinate systems
- **Git proficiency** – branching, PRs, rebasing, conflict resolution
- **Async programming** – experience with async/await patterns in Python
- **Fast-paced comfort** – you thrive in environments where requirements evolve based on user feedback
- **Basic frontend capability** – can read and write TypeScript/React when needed

Strongly Preferred

- Experience with **FastAPI** or similar async Python frameworks
- Familiarity with **PyTorch**, **YOLO**, or **OpenCV** for computer vision tasks
- Knowledge of **PDF processing** and text extraction techniques
- Understanding of **Celery** and Redis for distributed task processing
- Experience with **PostgreSQL** and SQLAlchemy (async)
- **AWS experience (S3, ECS, Textract)** and deploying infrastructure with Terraform
- Exposure to **LangChain**, **LangGraph**, or similar AI orchestration frameworks
- Experience containerizing applications with **Docker** for production deployments

Tech Stack Highlights

Backend: Python, FastAPI, Celery, PostgreSQL, Redis

AI/ML: OpenAI/Claude APIs, PyTorch, YOLO, OpenCV

Frontend: TypeScript, React, Next.js, Tailwind CSS

Infrastructure: AWS (S3, ECS, Textract), Docker, Terraform

What Makes This Role Different

- **Real impact:** Your work directly eliminates hours of manual labor for construction professionals
- **Customer-driven:** Weekly calls with users mean you see exactly how your features help (or don't)
- **Deep technical work:** Focus on challenging problems in document processing, computer vision, and AI integration
- **Learning-intensive:** Work with cutting-edge AI tools while solving complex real-world problems
- **Small team:** Your voice matters in technical decisions from day one
- **Collaborative environment:** Weekly in-person sessions for deep work and knowledge sharing
- **Deep listening:** AnchorGrid strives to maintain a culture and practice of deep listening

What We Offer

- The chance to join a vertical AI company on the ground floor
- The opportunity to impact the \$2.1T construction industry with an early stage team
- A competitive compensation package including salary and early stage equity

To Apply: Send your resume and a brief note about a technical problem you've solved involving document processing, AI, or computer vision to [careers@getanchorgrid.com]. Include links to relevant GitHub repos or portfolio projects.

AnchorGrid is an equal opportunity employer. We value diversity and are committed to creating an inclusive environment for all employees.