



GEMINUS

720 University Avenue Suite 200
Palo Alto CA94301 U.S.A.
Ph: (650) 234-0400

<https://geminus.ai/>
Twitter: @Geminus_AI
Email: info@geminus.ai

Full Stack Engineer

About Geminus

Geminus is the world's first *industrial process centric digital twin product* that delivers self-optimized design, predictive operational intelligence and asset performance enhancement to operators and engineering service providers. The product leverages *physics-constrained artificial intelligence* to deliver high-fidelity O&M intelligence and to continuously optimize the designs of assets & processes. Geminus integrates with operational data through SCADA systems, historian software, sensor gateways and enterprise asset management to trigger physics-constrained AI modeling of the end-to-end process for every significant time slice of data. This generates high-fidelity operational intelligence and design recommendations to prevent multi-step failures, adverse shifts in yield, structural process faults and degraded remaining useful lifetime.

The company is based in Palo Alto and venture funded by The Hive. It's co-founder and Chief Scientist, *Karthik Duraisamy*, has over a decade of rich experience in physics-based modeling and machine learning through both fundamental research as faculty at Stanford University & the University of Michigan and industry collaborations. The fast-growing R&D team combines experiences across physics modeling, AI, industrial domains and data operations. The product is in pilot deployments in oil & gas and hi-tech manufacturing verticals. The product is also being made available as a module in leading engineering simulation platforms.

About the Role

Geminus' product derives intelligence out of data from operational, experimental and numerical sources, and uses the inferences made to construct digital twins for industrial processes. This digital twin model development involves iterative amalgamation of machine learning and computational physics. The role of the Full Stack Engineer will be developing data processing components for the collection, curation and analysis of data. The role also includes building the data operations that binds the computational physics with the AI-driven training. The role leads the development of the core data architecture of the product.

About You

The successful candidate will have a proven track record of building complex data processing infrastructure:

- Degree in Computer Science, Engineering, or related fields
- Strong expertise in data structures, algorithms and coding
- Strong experience in development of data processing and online optimization systems in Python
- Experience with model training, analytics and serving with TensorFlow
- Experience with web infrastructure, data analytics, and distributed systems
- Disciplined approach to testing and quality assurance
- Experience with software engineering best practices (e.g. unit testing, code reviews, design documentation)
- Experience with AWS (e.g. IAM, EBS/S3, ECS, Lambda) or similar cloud computing environments.

Please send your resume to jobs@geminus.ai.

Geminus is an affirmative action employer and welcomes candidates who will contribute to the diversity of our team.