

IVAN LEON

Engineer, Researcher, Leader

☎ (803) 627-9500

✉ i.leonmaldonado@gmail.com

🌐 ielm.io

🐙 github.com/ielm

SUMMARY

I love diving into complex technical problems and building high impact teams and software solutions that are customized to solving them. In my role as Principal Engineer at Guidance Analytics, I successfully launched products ranging from AI/NLP healthcare solutions generating over \$15 million in value, to beautiful user interfaces that combine deep, complex information displays with state-of-the-art tools. I am - first and foremost - a systems thinker that applies bespoke strategy, management and people solutions synthesized from my deep understanding of lean methods and empathetic management. Above all, I believe in pursuing the elegant integration of technology in our everyday actions and interactions. My dream is to dismantle the barriers between technology and the human experience through thoughtfully designed agent systems and novel computing environments. I am deeply invested in designing the future of how we live, work, and play.

PROJECTS

Cargo SAM (Not Yet Public)

<https://github.com/cargo-sam>

A Cargo Subcommand for configuring and generating AWS SAM projects.

UploadThing API

<https://github.com/ielm/utapi-rs>

A high-level, ergonomic Rust crate for interacting with the Uploadthing API.

RELEVANT EXPERIENCE

Co-Founder

Monoql, Inc. | 08/2023 - Present

- Architected a distributed serverless infrastructure to optimize cost and scalability using AWS SAM, Lambda, Cognito, DynamoDB, and API Gateway.
- Established a pipeline for extracting clinical events from text and linking them to knowledge graph entities to streamline the data aggregation process for healthcare analytics.
- Developed RAG Agents for Question Answering on knowledge graphs and individual patient histories to enhance the accuracy of clinical decision-making and patient care management.
- Fine-tuned Large Language Models (LLM) for Named Entity Recognition (NER) and coreference resolution in clinical text, resulting in improved data extraction and analysis accuracy for clinical research and reports.
- Fine-tuned LLMs for personalized summarization of patient histories, providing tailored patient care insights and improving healthcare provider efficiency.

Rust / Python / SQL/NoSQL / AWS / HIPAA / HL7/FHIR / Serverless / Distributed Systems

Research Software Engineer - David Lab

Oregon State University | 10/2023 - 04/2023

- Engineered a data pipeline for extracting descriptive properties of individual taxa and microbial communities, aiding in advanced biological research and development.
- Developed a pipeline for extracting and integrating protein and microbial properties into knowledge graphs, facilitating enhanced data interconnectivity and research efficiency.
- Fine-tuned LLMs for learning microbiome properties using embedding techniques, advancing the lab's capabilities in predictive microbiology and microbial genomics.

Python / C/C++ / Fortran / AWS / LLMs / SQL

Principal Engineer & Director of Engineering

Guidance Analytics | 06/2022 - 11/2023

- Designed, implemented, and deployed rust-based services to process CMS MIPS measures and supply the backend architecture for a user-facing platform using GraphQL, gRPC, Docker, and Kubernetes
- Added "9s" of availability to the cloud-based services.
- Diagnosed the root cause of and removed long-tail latencies in the 99th percentile, resolving problematic delays across the distributed system infrastructure.
- Led efforts in optimizing infrastructure, achieving a 40% cost reduction and ensuring system scalability and reliability
- Developed large-scale data pipelines, ETL processes, and AI/ML Solutions that generated over \$15 million in value.
- Worked directly with administrators and IT of radiology groups across the country to integrate their EHR and billing systems into our data environment.

Rust / Python / AWS / Aurora / HIPAA / HL7/FHIR / GraphQL / gRPC / Docker

Software Engineering Manager

Guidance Analytics | 03/2022 - 06/2022

- Recruited and mentored a team of data scientists and engineers for the development of ML models and data pipelines, automating the evaluation of CMS MIPS measures and enhancing operational efficiency.
- Managed day-to-day operations, overseeing the product development lifecycle, ensuring timely delivery and quality.

Rust / Python / AWS / Docker / HIPAA / HL7/FHIR / Leadership / Radical Candor

AI Engineer

Guidance Analytics | 01/2022 - 06/2022

- Designed, built, and deployed information retrieval and classification systems to measure relevance and adherence to CMS MIPS Measures, improving measure evaluation and analytical accuracy
- Fine-tuned and optimized LLMs for classification and question answering to boost the system's predictive accuracy and end-user guidance.

Python / PyTorch / Pandas / Numpy / AWS / Docker / HL7/FHIR

Data Scientist

Meta | 04/2022 - 10/2022

- Developed user (A/B) experiment for Messenger Notifications, collecting data from 100+ million daily users.
- Engineered and optimized data processing and ETL pipelines that handle billions of data points daily.
- Worked on prediction model to determine whether to show notifications on Messenger or Blue based on user features.

Python / PyTorch / Scuba / Hive / Presto / DataSwarm (AirFlow) / A/B Testing

Founder & Software Engineer/Consultant

Omicron AI, LLC | 06/2018 - 01/2022

- Created the Omicron Framework for the development of knowledge-driven intelligent agents, OmniGraph knowledge representation and reasoning system, and the Delta Protocol for the sharing and access of data among networked agents.
- Provided consulting services on intelligent agents and knowledge graphs, leading to improved client capabilities in data analysis and decision-making.
- Implemented the Omicron Framework across various companies, enhancing their capabilities in developing intelligent agents and knowledge graphs.

Python / C/C++ / Knowledge Graphs / Sales / Consulting

Applied Science Researcher

WeWork | 04/2019 - 09/2019

- Developed a software system to automate space layouts with generative geometry algorithms, impacting 20+ million square feet of building space. This algorithm became one of the core pillars of Hypar.
- Worked on WeK8s, an internal tool for the orchestration of load-balanced services across the organization using Docker and Kubernetes.

GoLang / C/C++ / Generative Algorithms / APIs / Docker / Kubernetes

Research Scientist - LEIA Lab

Rensselaer Polytechnic Institute | 08/2018 - 01/2023

- Received a competitive three-year \$150k fellowship (RPI HASS Fellowship) to complete research in AI to incorporate Natural Language Understanding and Generation into the OntoAgent Cognitive Architecture.
- Led 3 end-to-end research projects, resulting in research papers that received distinguished research awards.
- Presented state-of-the-art neurosymbolic AI research that received an honorable mention at ACS2020.
- Designed and led implementation of a novel knowledge representation system for LLM-driven dialogue agents.
- Developed data models and microservices for CRUD operations accessible by cognitive agents for ontological and lexical representations with PostgreSQL & SQLite, Python, and GoLang.

Python / GoLang / PyTorch / Transformers / DSL / Ontology Engineering / SQL

Research Scientist - LACAI Lab

Rensselaer Polytechnic Institute | 08/2018 - 01/2023

- Researched behavioral modeling in social media and gaming environments, contributing to advanced understanding in human-computer interactions.
- Developed information extraction and retrieval methods for agent systems funded by the US Air Force and DARPA.
- Created a situationally-aware intelligent agent system, improving collaboration efficiency between humans and AI in achieving conversational goals.
- Developed a knowledge representation and reasoning system combining logic and probability, advancing intelligent agent decision-making processes.

Python / PyTorch / Transformers / Information Retrieval / Behavioral Modeling

EDUCATION

Ph.D - ABD

Cognitive | Computer Science

Rensselaer Polytechnic Institute

Master of Science

Cognitive Science

Rensselaer Polytechnic Institute

Bachelor of Architecture

Architecture

Rensselaer Polytechnic Institute

High School - Music Performance

French Horn

SCGSAH

SKILLS

Languages: Rust, Golang, Python, Elixir, C/C++, JavaScript/Typescript, SQL, Java.

Technologies: AWS, Docker, Kubernetes, PostgreSQL, Hive, Elasticsearch, Tantivy, Solr/Lucene.

Management Skills: Lean Methodology, Radical Candor, Empathetic management, Cross-functional Communication, Active Listening.

Languages: English - Native, Spanish - Native