

# SEPEHR SADIGHPOUR

[hi@sepehr.xyz](mailto:hi@sepehr.xyz)

(510) 759-4850

## AI PLATFORM ARCHITECT

---

### Profile

Senior Staff ML engineer

---

### Experience

#### Grindr

##### Senior Staff MLOps Engineer

Feb 2024 - Present

Vendor evaluation and cost management, in particular around LLM usage, with high focus on privacy, security, and regulatory compliance.

Establishing best-practices for end-to-end AI projects and leading cross-team collaboration.

Advising on LLM use-cases, risks, and costs, across a variety of applications.

Deploying chatbots that can scale to serve millions of requests per day on AWS EKS.

Implemented image detection pipelines capable of processing millions of images a day.

*Databricks, Snowflake/Cortex, AWS SageMaker, AWS Bedrock, AWS EKS, Terraform, LLM evaluation platforms*

---

#### Equilibrium Energy

##### Staff ML Engineer

Oct 2022 - Feb 2024

Designed and built a new ingestion framework based on Temporal.io and dbt. Reduced duration of backfills from days to minutes, while reducing cost and improving resilience to outside outages

Implemented distributed hyper-parameter tuning of Deep Learning models in Metaflow

Tracked DORA and ML Test Score for communicating ML infra maturity to management

Simplified Feast feature store management to allow Data Scientists to self-serve

Developed dashboard for visualizing time-series forecast output and model performance

Introduced Great Expectations to validate data at ingestion

Designed experimentation workflow for data scientists, allowing them to self-serve via sandboxed clones of data warehouse & personal dbt development workflow

*Technologies: Terraform, Docker, Kubernetes, AWS EKS, Metaflow, Temporal.io, dbt, Snowflake (Streams & Tasks for CDC), Airflow MWAA, Great Expectations, pants, Feast, Argo Workflows, Argo CD, Tensorflow, XGBoost, Optuna, Dash*

---

#### Actium Health

##### Staff Data Engineer & Manager

Mar 2021 - Oct 2022

Managed and mentored team of 4 data engineers, acting as tech-lead and scrum master.

Introduced rigorous standards of validating data pipelines and monitoring system behavior.

Identified and addressed gaps in visibility, monitoring, and alerting

*Technologies: Jira, Kubernetes, prefect, Great Expectations*

---

#### Bodyport

##### Senior Data Engineer

May 2020 - Mar 2021

Built tooling for visualizing and annotating physiological time-series data from device for use in machine learning applications

Set up production replicas, ETL, and data warehouse for analytical and research purposes

Brought all data-science related infrastructure under code, creating a completely reproducible tech stack.

Created developer tools for generating realistic data while protecting patient privacy.

*Technologies: Dash, Terraform, Azure App Service, postgres+pglogical, scipy, scikit-learn*

---

## Clarify Health Solutions

### Senior ML Engineer

June 2018 - Mar 2020

Architected ML training and tracking framework capable of parallel training and serving 1M+ models in production in batch and real-time modes.

Implemented distributed model interpretation framework in Spark.

Built Spark pipelines for data validation, feature engineering, and model training on 10s of billions of rows of healthcare claims and social determinants of health data.

### ML Engineer

Nov 2017 - June 2018

Trained risk models for a variety of health outcomes, implementing requirements from statistics experts and translating government policy documents (Medicare/Medicaid)

*Technologies: H2O, AWS (ECS/Batch/EMR/Redshift), Spark+Spark ML, dagster*

---

## Omada Health

### Data Scientist

Apr 2016 - Jul 2017

Guided product direction via analyses of trends in participant behavior, leading to 6% lift in user engagement

Designed and implemented A/B tests for new features in close collaboration with product

Managed 4-person data team's research into areas such as chatbots, reinforcement learning, and machine vision

Proposed and prototyped multi-armed bandit for more data-efficient personalization

Used Word2Vec & Doc2Vec to cluster patient communications.

*Technologies: pandas, gensim, VowpalWabbit, R/Shiny/ggplot*

---

## Galvanize University

### Data Science Fellow

Oct 2015 - Jan 2016

Capstone Project: Generated text embeddings of scientific publications and built an information retrieval app for searching Arxiv semantically, using vector length to control temperature of results.

---

## BayThrive

### Owner, Developer, Project Manager

Jan 2010 - Apr 2016

Increased traffic 20x and conversion rates 4x for biotech sales firm through SEO and iterative design.

Built development and production infrastructure for early-stage AI startup using Docker and Ansible.

Developed heuristic for financial research firm to estimate degree of self-similarity in time-series data.

---

## Mepkin Abbey

### Webmaster, Go-To Guy

Jun 2008 - Jan 2010

Built WordPress website, e-commerce solution, and trained store employees in inventory management and fulfilling online sales.

Managed development of custom registration system, from need-assessment to hiring and project management.

---

## Duke University

Bachelors of Science, Physics

Class of '08

Capstone Project: Trained neural networks to solve for energy states of quantum systems by representing the system Hamiltonian. Used genetic algorithms to reduce training time.