

Zachary Tessler, Ph.D.

New York, NY | 646-727-0667 | ztessler@gmail.com

Experienced DS/ML leader with technical expertise in machine-learning and data/ML pipelines; bio/environmental/physical system modeling using deep-learning, probabilistic Bayesian inference, and simulation; and IoT/time-series forecasting. Coaching and leading teams working across the data engineering, data science, and MLEng/MLops spectrum. Experienced in dynamic and fast-paced startup environments.

Dec. 2023 - Present:

Danelec Marine, Denmark (remote from NY), *Director of Data Science and Engineering*

Lead Engineering and Data Science teams comprising Danelec's US-based R&D organization, focused on ship performance and voyage optimization products and technologies.

- Recruited and retained top-performing SWE and DS staff following Danelec acquisition of Nautilus Labs.
- Designed and implemented post-acquisition engineering organization structure.
- Led definition and oversaw execution of high-impact but low-dependency engineering projects (incl. IoT time-series database migration, and service orchestration migration from ECS to k8s) for performance, scaling, and cost improvements. Initial results show 80% reduction in typical request times through improved compute/storage trade-offs.
- Coordinated and aligned engineering and data-science strategy with product and client-team goals.

Jan. 2022 - Dec. 2023:

Nautilus Labs, New York, NY, *Director; Data Science; Data Science Manager*

Led team responsible for design, implementations, and management of all data science and machine learning tools, products, and services for Nautilus Labs, an innovative maritime data and fuel consumption optimization startup.

- Represented DS and ML function in all company-wide strategic planning; led DS strategy and research pipeline; technical owner of all ML model design/training/deployment; primary DS communicator for clients and investors.
- Grew NL DS/ML function from 1 to 6 team members. Led hiring process for multiple junior and multiple senior data scientists, and managed internal transition of a machine learning engineer onto the DS team. Increased team velocity 66% in second half of 2023.
- Proposed, prototyped, and managed migration of legacy model-training batch job to a modern, parallelized, orchestrated (Dagster) ML pipeline on AWS EKS (Kubernetes), improving scalability, configurability, observability and testing, while decreasing total job runtime by over 7x, with no change to cost. IaC via Terraform, CI/CD via Github Actions.

Sept. 2019 - Dec. 2021:

Nautilus Labs, New York, NY, *DS Tech Lead; Senior Data Scientist; Data Scientist*

Led research, development, and implementation of machine learning solutions for Nautilus Labs. Focused on ship performance modeling and prediction, weather risk assessment and monitoring, IoT data reliability, and outlier detection.

- Developed state-of-the-art probabilistic Bayesian modeling solution for clients without IoT sensors, unlocking 62% of the expected accuracy improvement from IoT sensor installation, but with 0 new CapEx requirement

- Re-designed and implemented improved platform-wide vessel performance modeling strategy, adding flexibility to support inconsistent client data availability, and increasing client coverage from 69% to 94%.
- Developed a semi-automated modeling library (Python) to codify domain-based modeling best-practices, manage model deployments, persist serialized training configuration, and automate model lifecycle.

Apr. 2012 - Sept. 2019:

CCNY and Advanced Science Research Center, City University of New York, New York, NY

Director of the Coastal and Ocean Science Synthesis Facility and Research Assistant Professor; Postdoctoral Researcher

Developed and led a research program into coastal environmental change, flood risk, and vulnerability. Synthesized observational, remote sensing, and numerical model data sources using analytical and machine-learning models.

- Led team awarded \$1.7M, 3-year NASA grant to develop predictive geospatial models of coastal ocean hypoxia risk by integrating remote sensing (*e.g.*, MODIS, MERIS, LandSat, Sentinel-1/2/3) and coupled met-ocean numerical modeling.
- Developed a predictive forecast model of future sea-level rise and anthropogenic pressures in deltas using end-to-end data and modeling pipelines. Published in *Geomorphology* and *Sustainability Science* journals.
- Lead author of highly-cited (top 5% of all research outputs, per Altmetric) research article in *Science* on contemporary and future coastal flood risk and vulnerability due to sea level rise and land subsidence in 48 global deltas.

Key Projects:

- CI/CD-based model training pipeline for improved data scientist efficiency, collaboration, reproducibility, and leveraging of cloud resources. *Technologies:* Dagster, Github Actions, Terraform, Kubernetes, Python, GraphQL. *Role:* Individual Contributor, Manager/Director
- Refactoring of request-time calculation as pre-computed and persisted data for improved performance and scaling. *Technologies:* Postgresql, Golang, AWS Lambda. *Role:* Manager/Director
- Research, development, and productionalization of Bayesian Hierarchical Models for transfer-learning between entities with- and without IoT sensor-enabled systems. *Technologies:* Python, PyMC, Git, Dagster, Kubernetes. *Role:* Individual Contributor, Manager/Director

Education:

Ph.D., Physical Oceanography (Ocean Physics), Columbia University, New York, NY, 2012

M. Education, Curriculum and Instruction, Arizona State University, Tempe, AZ, 2006

B.S., Geological Science, Geo/Bio Concentration, Brown University, Providence, RI, 2004

Technical Skills: *Languages:* Python, SQL, MATLAB, C, some Golang; *Frameworks/Libraries/APIs:* Numpy, Pandas, Scikit-learn, XGBoost, PyTorch, PyMC; *Tooling:* Git/GitHub, Linux/OSX/Win, Docker, AWS EKS/ECS/Fargate/Sagemaker, Dagster, MLFlow, Terraform, Github Actions, CircleCI