Navdeep Gill

mr.navdeepgill@gmail.com github.com/navdeep-G Google Scholar
linkedin.com/in/navdeep-gill-b1729456

Education

California State University, East Bay, Hayward, CA

M.S. in Statistics (Emphasis on Computational Statistics) *Sep 2012 – Jun 2014*

- *Recipient of the NSF Statistics Scholarship (2012–2013)*
- *Teaching Assistant for Introductory Statistics (Instructor: Dr. Lynn Eudey)*

B.A. in Psychology; B.S. in Statistics (Minor: Mathematics) *Aug 2010 – May 2012*

- *Focused coursework on statistical modeling, experimental psychology, and advanced mathematics.*
- *Engaged in research projects combining statistical analysis and psychological studies.*

Professional Experience

ServiceNow, Santa Clara, CA

Staff Senior Product Manager, Responsible AI *Aug 2024 – Present*

- Contribute to defining the vision, strategy, and roadmap for Responsible AI to ensure ServiceNow meets the highest standards of quality, fairness, accountability, and transparency in AI product development and delivery.
- Work within a cross-functional team of product managers, engineers, legal experts, risk specialists, and researchers to build ethical and trustworthy AI solutions.
- Develop and communicate the Responsible AI vision, strategy, and roadmap across ServiceNow and with external stakeholders.
- Define and monitor key performance indicators and success metrics for the Responsible AI & AI Governance program.
- Partner with engineering, design, legal, risk, and compliance teams to operationalize AI Governance & Responsible AI policies.
- Support the design and delivery of Responsible AI training and continuing education programs for employees and customers.
- Help establish relationships with key customers, partners, and thought leaders in the Responsible AI domain.

H2O.ai, Mountain View, CA

Lead Data Scientist *Aug 2021 – Aug 2024*

- Led engineering and data science initiatives to create innovative solutions in AI governance and responsible AI, architecting scalable systems using Python, Go, Java, TypeScript, Kubernetes, AWS, and Docker.
- Managed the software development lifecycle, enforcing best practices, coding standards, and quality assurance to ensure technical alignment and product integrity.
- Collaborated with product teams to define architecture, roadmaps, and features for AI governance products, aligning technical implementations with strategic goals.
- Engaged with key customers and stakeholders to translate requirements into actionable technical solutions, addressing customer needs while aligning with product strategy.

- Advised the executive team on AI governance strategy, providing technical insights and recommendations to shape business objectives.
- Contributed thought leadership through academic publications, white papers, and blogs, establishing the organization as a leader in responsible AI.
- Managed a global team of engineers and data scientists across three continents, fostering collaboration, driving innovation, and delivering complex projects.

Senior Software Engineer *Sep 2017 – Aug 2021*

- Spearheaded engineering and data science initiatives to enhance machine learning interpretability in Driverless.ai, H2O.ai's flagship AutoML platform.
- Designed and implemented scalable solutions using Python, Java, and TypeScript to build reliable and interpretable machine learning tools.
- Developed advanced model explanation tools, debugging frameworks, and fairness metrics to address bias and ensure responsible AI outcomes.
- Engineered interpretable models, including decision trees and generalized additive models (GAMs), increasing transparency and usability for end-users.
- Led UI development efforts with TypeScript to create intuitive interfaces, improving accessibility and adoption of interpretability features.
- Collaborated with product managers and cross-functional teams to align architecture and roadmaps with strategic goals, enhancing platform reliability.
- Engaged with customers to understand requirements and translate them into actionable technical solutions, focusing on fairness and compliance.
- Published research on model interpretability and fairness, establishing thought leadership and driving innovation in responsible AI.

Software Engineer *Oct 2015 – Sep 2017*

- Played a key role in advancing H2O AutoML by automating data preprocessing, model selection, hyperparameter tuning, and model evaluation processes.
- Contributed to the development of H2O4GPU, leveraging GPU capabilities to accelerate machine learning tasks such as model training, inference, and unsupervised learning.
- Optimized multiple machine learning algorithms for efficient GPU usage, significantly reducing execution times.
- Played a pivotal role in enhancing H2O-3, focusing on distributed machine learning algorithms and scalable data processing workflows.
- Developed the rsparkling package, integrating H2O with Spark's in-memory processing via the sparklyr API, enabling seamless workflows between Spark SQL and H2O.
- Delivered efficient, production-ready solutions for model training, deployment, and evaluation across large-scale datasets.

Cisco, San Jose, CA

Business Intelligence Engineer *Sep 2014 – Oct 2015*

- Played a key role within an engineering team dedicated to designing and building advanced analytics solutions for Cisco's call center platform.
- Spearheaded the development of efficient ETL (Extract, Transform, Load) processes utilizing Talend Integration Suite (TIS) to streamline data processing workflows.
- Utilized MS SQL and PostgreSQL to architect and implement database schemas optimized for in-depth analysis of call center data.

- Applied Python and R for exploratory data analysis, data cleaning, manipulation, and visualization tasks, preparing data for advanced analytics.
- Developed interactive dashboards and descriptive analytics using Python and Tableau, delivering actionable insights to technical and non-technical stakeholders.
- Designed and implemented machine learning models for call volume forecasting, sentiment analysis, anomaly detection, call routing optimization, and customer segmentation.
- Engineered a robust KPI (Key Performance Indicator) framework, including metrics like first call resolution rate, average handling time, Customer Satisfaction Score (CSAT), and agent utilization rate.
- Created essential metrics and analytical models to enhance customer effort and agent performance effectiveness.
- Presented key insights derived from analytics and machine learning models to stakeholders, ensuring clear communication and actionable takeaways.
- Maintained adherence to project timelines using Agile methodologies, ensuring consistent and timely deliveries.
- Collaborated on initiatives to improve system scalability, optimize query performance, and implement data governance best practices.

FICO, San Rafael, CA

Analytic Science Consultant *Apr 2013 – Sep 2014*

- Developed, validated, and implemented machine learning models across the credit risk management lifecycle, including originations, customer management, collections, and loss forecasting.
- Led end-to-end model development processes, encompassing data preprocessing, feature engineering, segmentation analysis, model building, and deployment.
- Utilized advanced statistical and machine learning techniques to create predictive models, improving risk assessment accuracy and financial decision-making.
- Conducted extensive data analysis to identify key trends, optimize model performance, and ensure adherence to industry standards and best practices.
- Collaborated cross-functionally with business, compliance, and technical teams to align models with regulatory requirements and client objectives.
- Delivered actionable insights through detailed reports and presentations, empowering clients to mitigate risk and enhance portfolio management strategies.

Academic Experience

University of California, San Francisco

Research Assistant *May 2012 – Apr 2013*

- Investigated neural mechanisms of memory and attention under the guidance of Dr. Bradley Voytek, focusing on changes in normal aging and dementia, with emphasis on therapeutic interventions.
- Leveraged real-time EEG tracking, non-invasive brain stimulation (tDCS), and behavioral measurements to study cognitive processes.
- Performed comprehensive data analysis using MATLAB and Python, extracting meaningful insights from complex datasets.
- Trained fellow research assistants in data collection, experiment administration, and data analysis techniques.

- Managed participant recruitment, consent, experimental sessions, and post-session debriefing/payment processes.
- Co-authored the article *"Preparatory Encoding of the Fine Scale of Human Spatial Attention"* in the *Journal of Cognitive Neuroscience*.
- Presented research findings at the Cognitive Neuroscience Society Meeting and Entertainment Software and Cognitive Neurotherapeutics Society Conference.

The Smith-Kettlewell Eye Research Institute

Research Assistant *May 2011 – May 2012*

- Studied 3D depth perception and eye movement control under the guidance of Dr. Christopher Tyler, focusing on the effects of brain injuries on visual processing.
- Employed eye-tracking and fMRI imaging techniques to investigate the role of various cues and eye movements in depth perception.
- Processed and analyzed psychophysical and fMRI datasets using MATLAB, Python, R, and MS Excel.
- Conducted data analysis to uncover patterns and insights from behavioral and neuroimaging studies.
- Managed participant recruitment, informed consent, experimental sessions, and post-session debriefing/payment processes.
- Co-authored articles: *"Hysteresis in Stereoscopic Surface Interpolation: A New Paradigm"* and *"Analysis of Human Vergence Dynamics"* in the *Journal of Vision*.

California State University, East Bay

Undergraduate Researcher *Jun 2009 – May 2011*

- Conducted research under Dr. David Fencsik on human visual processing and short-term memory mechanisms.
- Investigated visual memory processes with a focus on the extraction and temporary storage of environmental information.
- Analyzed behavioral data using MATLAB, Python, R, and MS Excel, extracting key insights from experimental results.
- Managed participant recruitment, consent procedures, experiment administration, and post-session debriefing/payment processes.
- Presented findings at undergraduate research conferences, including the UC Berkeley Cognitive Science Conference and Psychology Undergraduate Research Conference.

Teaching

Guest Lecturer

- **QST BA 840: Data Ethics** – Graduate Course, Boston University Questrom School of Business *Spring 2022*
- **FNCE 3490: Data Science and Business Analytics** – Graduate Course, Santa Clara University *Spring 2016*

Publications

- **Gill, N., Zhang, S.** (2023) *Guidelines for Effective AI Governance with Applications in H2O AI Cloud.* H2O.ai. URL

- **Gill, N., Mathur, A., Conde, M.** (2022) *A Brief Overview of AI Governance in Responsible Machine Learning Systems.* NeurIPS Workshop on Trustworthy and Socially Responsible Machine Learning (TSRML).
- **Hall, P., Gill, N., Cox, B.** (2020) *Responsible Machine Learning: Actionable Strategies for Mitigating Risk and Driving Adoption.* O'Reilly Media, Inc.
- **Gill, N., Hall, P., Montgomery, K., Schmidt, N.** (2020) *A Responsible Machine Learning Workflow with Focus on Interpretable Models, Post-hoc Explanation, and Discrimination Testing.* Information, 11(3):137.
- **Hall, P., Gill, N.** (2019) *An Introduction to Machine Learning Interpretability, Second Edition: An Applied Perspective on Fairness, Accountability, Transparency, and Explainable AI.* O'Reilly Media, Inc.
- **Hall, P., Gill, N., Schmidt, N.** (2019) *Proposed Guidelines for the Responsible Use of Explainable Machine Learning.* NeurIPS Workshop on Robust AI in Financial Services.
- **Hall, P., Gill, N., Meng, L.** (2018) *Testing Machine Learning Explanation Techniques.* O'Reilly Media.
- **Hall, P., Gill, N.** (2018) *An Introduction to Machine Learning Interpretability: An Applied Perspective on Fairness, Accountability, Transparency, and Explainable AI.* O'Reilly Media, Inc.
- **Hall, P., Gill, N., Kurka, M., Phan, W.** (2017) *Machine Learning Interpretability with H2O Driverless AI.* H2O.ai. URL
- **Voytek, B., Samaha, J., Rolle, C. E., Greenberg, Z., Gill, N., Porat, S.** (2017) *Preparatory Encoding of the Fine Scale of Human Spatial Attention.* Journal of Cognitive Neuroscience, 29, 1302–1310.
- **Tyler, C.W., Elsaid, A.M., Likova, L.T., Gill, N., Nicholas, S.C.** (2012) *Analysis of Human Vergence Dynamics.* Journal of Vision, 12(11): 21.

Conference Presentations

- **Gill, N., Montgomery, K.** (2024) *Interpretability for Generative AI.* H2O GenAI Day, Atlanta, GA, January 23.
- **Gill, N.** (2023) *Guardrails for LLMs.* H2O Open Source GenAI World, San Francisco, CA, November 7.
- **Gill, N., Mathur, A.** (2022) *Incorporating AI Governance to Increase Adoption in Business Applications.* MLOps World 2022, New York, NY, July 14.
- **Gill, N., Tanco, M.** (2021) *Security Audits for Machine Learning Attacks.* MLOps World 2021, June 16.
- **Gill, N.** (2021) *Training Understandable, Fair, Trustable and Accurate Predictive Modeling Systems.* Duke Machine Learning Day, Durham, NC, March 27.
- **Gill, N.** (2019) *Human Centered Machine Learning.* Artificial Intelligence Conference, San Jose, CA, September 11.
- **Gill, N.** (2019) *Interpretable Machine Learning Using rsparkling.* Symposium on Data Science and Statistics, Bellevue, WA, May 31.
- **Gill, N.** (2019) *Practical Machine Learning Interpretability Techniques.* GPU Technology Conference, San Jose, CA, March 21.
- **Gill, N.** (2018) *Distributed Machine Learning with H2O.* Joint Statistical Meeting, Vancouver, Canada, August 1.
- **Gill, N.** (2018) *H2O AutoML.* Symposium on Data Science and Statistics, Reston, VA, May 16.
- **Hall, P., Gill, N., Chan, M.** (2018) *Practical Techniques for Interpreting Machine Learning Models: Introductory Open Source Examples using Python, H2O and XGBoost.* 1st ACM Conference on Fairness, Accountability, and Transparency, New York City, NY, February 23–24.
- **Gill, N., Hall, P., Chan, M.** (2017) *Driverless AI Hands-On Focused on Machine Learning Interpretability.* H2O World, Mountain View, CA, December 11.
- **Gill, N.** (2017) *From R Script to Production using rsparkling.* Spark Summit, San Francisco, CA, June 14.
- **Gill, N.** (2016) *Scalable Machine Learning in R with H2O.* useR Conference, Stanford, Palo Alto, CA, July 11.
- **Voytek, B., Porat, S., Chamberlain, J., Balthazor, J., Greenberg, Z., Gill, N., Gazzaley, A.** (2013) *Examining the Efficacy of the iPad and Xbox Kinect for Cognitive Science Research.* 2nd

Annual Meeting of the Entertainment Software and Cognitive Neurotherapeutics Society, Los Angeles, CA, March 15–17.

- **Greenberg, Z., Gill, N., Porat, S., Samaha, J., Kader, T., Voytek, B., Gazzaley, A.** (2013) *Increased Visual Cortical Noise Decreases Cued Visual Attention Distribution.* 20th Annual Meeting of the Cognitive Neuroscience Society, San Francisco, CA, April 13–16.
- **Tyler, C.W., Gill, N., Nicholas, S.** (2012) *Hysteresis in Stereoscopic Surface Interpolation: A New Paradigm.* 12th Annual Meeting of the Vision Sciences Society, Naples, FL, May 11–16.
- **Gill, N., Fencsik, D.** (2012) *Effects of Disruptions on Multiple Object Tracking.* California Cognitive Science Conference, UC Berkeley, CA, April 28.
- **Gill, N., Fencsik, D.** (2011) *Effects of Distractions on Recovery Time.* Psychology Undergraduate Research Conference, UC Berkeley, CA, May 1.

Patents

- **Chan, M., Gill, N., Hall, P.** (2024) *Model Interpretation.* U.S. Patent No. 11,922,283. Washington, DC: U.S. Patent and Trademark Office.
- **Chan, M., Gill, N., Hall, P.** (2022) *Model Interpretation.* U.S. Patent No. 11,386,342. Washington, DC: U.S. Patent and Trademark Office.