Dheeraj Dhillon

ddhillon@wisc.edu | +1(608)721-4309 | linkedin.com/in/djdhillxn | djdhillxn.github.io

EDUCATION

University of Wisconsin-Madison **Aug 2025 - May 2027**
Master of Science (M.S.), Computer Science
Coursework: Reinforcement Learning, Machine Learning, Game Theory, Operating Systems

Indian Institute of Technology (IIT) Roorkee **Jul 2019 - May 2023**
Bachelor of Technology (B.Tech.), Electronics and Communication Engineering Cumulative GPA: 8.491/10.0

SKILLS

Programming Languages: Python, C++, C, Spark, SQL
Frameworks: PyTorch, NumPy, Gymnasium, ALE, MuJoCo, Git, transformers, scikit-learn, statsmodels, spaCy, NLTK

WORK EXPERIENCE

Gartner Inc., Associate Data Scientist, Gurgaon **Mar 2024 - Aug 2025**

- Built and productionized a 3-level topic modeling pipeline in Databricks/Spark over 20k user queries and 30k Gartner research documents using all-MiniLM-L6-v2 embeddings and Louvain clustering to surface demand-sensing insights.
- Developed Implied Value Rating (IVR), a document-level metric estimating the uplift in user retention from readership, used to inform search reranking toward high-impact content.
- Modeled IVR for 21k+ documents with Bayesian logistic regression using engagement data from 100k licensed users and priors derived from churn models.
- Explained drivers of document performance by regressing IVR against topic profiles, uniqueness, and engineered semantic, discourse, and NER features, identifying 50+ significant attributes and achieving adjusted R2 = 0.62 on 4k IT End User documents.

HiLabs Inc., Data Scientist, Bangalore **Jul 2023 - Mar 2024**

- Built an information extraction pipeline for 20k Electronic Medical Records (EMRs) to extract structured Subject-Predicate-Object (SPO) triplet relationships using lexico-syntactic hyponym patterns.
- Developed a Named Entity Recognition (NER) model with linear-chain Conditional Random Fields (CRF) architecture for BIO tagging covering 30 clinical entity types (disease, dosage, drug) reaching a span F1-score of 82% on EMRs.
- Evaluated contextual and linear bandit strategies for personalized Warfarin dosage selection under sequential feedback, comparing Epsilon-Greedy, Upper Confidence Bound (UCB), Thompson Sampling, and LinUCB.
- Benchmarked online decision accuracy at 68% on dose selection, highlighting the trade-off between exploration and patient-specific treatment optimization.

Microsoft R&D, Data and Applied Scientist Intern, Hyderabad **May - Jul 2022**

- Deployed a BM25-based message recommender for Microsoft Teams Search. Integrated Latent Dirichlet Allocation (LDA) topic-aware query expansion increasing NDCG@10 to 40% from 31% on a new dataset of 8,000 messages.
- Implemented TextRank and KeyBERT algorithms for keyphrase extraction (KPE). Formulated a sentence-transformers embedding-based evaluator with an adjustable cosine similarity threshold for semantic matching.
- Achieved optimal results for the KPE task on the Teams dataset using KeyBERT with an F1@5 of 55% using cosine similarity metric, up from 43% obtained using hard-matching.

PUBLICATIONS

InfraNet: An Ensemble Approach for Real-time Wildlife Detection using Infrared Thermal Imaging, IEEE AVSS **Aug 2025**

PROJECTS

Trust Region Policy Optimization and Policy Gradient Methods, UW-Madison **Mar-Apr 2026**

- Implemented Trust Region Policy Optimization (TRPO) from scratch in PyTorch, from trajectory collection and advantage estimation to KL-constrained policy optimization using conjugate gradient updates, Fisher-vector products, and line search.
- Implemented PPO with both clipped and KL-penalized objectives, and benchmarked TRPO, PPO, and NPG on 3 MuJoCo locomotion tasks and 7 Atari games using Gaussian MLP policies for control and categorical CNN policies for Atari.
- Analyzed monotonic policy improvement guarantees in CPI and TRPO through surrogate performance objectives, clarifying how mixture-update constraints in CPI and KL-divergence constraints in TRPO stabilize policy updates.

Word-Guessing Simulator **Feb 2026**

- Designed a word-guessing simulator with a greedy letter-selection policy using bidirectional character n-grams (n=6 to 2), with Katz backoff interpolation, Kneser-Ney smoothing, and start-end word padding to improve robustness under sparse contexts.
- Attained a 64% win rate on a test set of 170k dictionary words with a 6 mistake budget, trained on a set of 227k words.