

Nikhil Sheoran

✉ nikhilsheoran96@gmail.com
in [linkedin.com/in/nikhil96sher](https://www.linkedin.com/in/nikhil96sher)

☎ (217) 318-9259
📄 github.com/nikhil96sher

EDUCATION

University of Illinois at Urbana Champaign

Master of Science in Computer Science

Thesis: DeepOLA - An Online Aggregation Approach to Approximating Deeply Nested Queries

GPA: **3.95/4.0** | **Advisor:** Prof. Yongjoo Park

Champaign, IL
Aug 2021 - Dec 2022

Indian Institute of Technology Roorkee

Bachelor of Technology in Computer Science and Engineering

GPA: **9.52/10.0** | **Advisor:** Prof. Sugata Gangopadhyay and Prof. Manoj Mishra

Roorkee, India
June 2014 - May 2018

EXPERIENCE

Databricks

Software Engineer, Query Optimization

Mountain View, CA
Feb 2023 - Present

Meta Platforms, Inc. (formerly Facebook)

Software Engineer Intern, AI Productivity

Menlo Park, CA
May 2022 - Aug 2022

- Ranking models suffer from normalized entropy explosion (NEX) when the model performance gets poorer than a static baseline.
- Developed a data-driven classification model to predict NEX whose attribution scores can be used to perform root-cause-analysis.
- Utilized input features' statistics and model architecture representation obtained using unsupervised embeddings.

Adobe Research

Research Associate, Big Data Experience Lab

Bangalore, India
June 2018 - Aug 2021

- *Predicate-Aware Approximate Query Processing*
 - Built a conditional generative model that uses query's predicates to generate targeted samples.
 - Evaluated query approximation error, latency and memory footprint against multiple baselines.
- *Scheduling of Time-Varying Workloads in Multi-Tenant Clusters*
 - Built a Deep-RL based agent for taking scheduling decisions exploiting temporal resource usage patterns.
 - Evaluated the performance with average resource-utilization, fragmentation and over-utilization on real production traces.
- *Multi-Touch Attribution for B2B Marketing Journeys*
 - Built a conversion prediction model with each stage transition modelled using a temporal convolution network.
 - Utilized layer-wise relevance propagation for computing attribution scores validated with perturbation-based techniques.
- *Surveys Without Questions: A Reinforcement Learning Approach*
 - Modelled users' behavior on an online platform as a partially-observed MDP and extracted proxy ratings from clickstream.
 - Evaluated the derived user experience metric against survey scores and provided aggregate-level actionable insights.

Adobe Research

Research Intern, Big Data Experience Lab

Bangalore, India
May 2017 - Jul 2017

- Modelled the temporal nature of users' online browsing behavior through various models - constrained LSTM, Probabilistic Suffix Tree and Hidden Markov Models.
- Proposed the concept of stage-wise experience values and their computation based on user's behavior logs.

PROJECTS

DeepOLA: Online Aggregation for Deeply Nested Queries [\[Abstract\]](#) [\[Thesis\]](#) [\[Repository\]](#)

Masters Thesis Project, CreateLab

UIUC
Jan 2022 - Present

- Propose a generalized framework and implementation for performing online aggregation in arbitrarily nested queries.
- Introduce the notion of *incremental dataframes* and *incremental operations* to support online aggregation.

Analysis of RCA Tools for Cloud Applications [\[Report\]](#)

Course Project, Reliability of Cloud Systems

UIUC
May 2022

- Automated RCA tools help identify potential root causes in cloud applications given an alert.
- Compared two tools - *Sieve* and *MicroRCA* in their performance in detecting root causes after manually injecting various faults.

Distributed Storage Networks with Smart Contracts Incentivisation [\[Report\]](#)

Bachelor Thesis Project, Department of Computer Science

IIT Roorkee
Jan 2018 - April 2018

- Proposed a smart-contract based storage network incentivized for sharing storage.
- Obtains proof of space (availability of storage) through memory-hard puzzles based on Merkle trees.

SKILLS

Languages: Python, Rust, Scala, C++, SQL, Bash, Java, PHP

Frameworks/Utilities: Apache Spark, Keras, PyTorch, Tensorflow, Django, HTML, CSS, Nginx.

HONORS

- AIR 10 in ACM ICPC Chennai On-Site Regionals 2017-18.
- AIR 5 in Microsoft Build The Shield Onsite Round 2016.
- Prime Minister's Scholarship Scheme Award 2014-18.

- KVPY Fellowship Award 2013.

OTHER ACTIVITIES

- Graduate Teaching Assistant, Fall 2022. CS511: Advanced Data Management, UIUC.
- Graduate Research Assistant, Spring 2022. Prof. Yongjoo Park, UIUC.
- Graduate Teaching Assistant, Fall 2021. CS240: Introduction to Computer Systems, UIUC.
- Chief Coordinator, Information Management Group, IIT Roorkee
- Vice Chair, ACM Student Chapter, IIT Roorkee
- Mentor, Student Mentorship Programme, IIT Roorkee

PUBLICATIONS (8)

A Step Towards Deep Online Aggregation

[\[Paper\]](#)

N. Sheoran, S. Chockchowwat, A. Chheda, S. Wang, R. Verma, Y. Park
In Proceedings of the ACM on Management of Data, SIGMOD 2023.

Electra: Conditional Generative Model Based Predicate-Aware Query Approximation

[\[Paper\]](#)

N. Sheoran, S. Mitra, S. Ghetia, J. Varshney, V. Porwal, T. Mai, A. Rao, V. Madukkuri, L. Mishra.
In Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022.

Scheduling of Time-Varying Workloads in Multi-Tenant Clusters using Deep Reinforcement Learning

[\[Paper\]](#)

S. Shanka*, N. Sheoran¹, S. Mitra.

In Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021.

Learning to Place Applications in a Shared Cluster

[\[Slides\]](#) [\[Paper\]](#)

S. Mitra, S. Shanka, N. Sheoran, N. Dhake, R. Nehra, R. Simha.

In Proceedings of the 10th ACM SIGOPS Asia-Pacific Workshop on Systems, APSys 2019.

Efficient Insights Discovery through Conditional Generative Model based Query Approximation

[\[Talk\]](#) [\[Paper\]](#)

V. Porwal, S. Mitra, F. Du, J. Anderson, N. Sheoran, A. Rao, T. Mai, G. Kowshik, S. Nair, S. Arora, S. Mahapatra
In Proceedings of the International Conference on Management of Data - Demo Track, SIGMOD 2022.

Multi-touch Attribution for Complex B2B Customer Journeys using Temporal Convolutional Networks

[\[Paper\]](#)

A. Agrawal, N. Sheoran, S. Suman and G. Sinha.

In the Companion Proceedings of TheWebConf 2022.

Surveys Without Questions: A Reinforcement Learning Approach

[\[Paper\]](#)

A. Sinha, D. Jain, N. Sheoran, S. Khosla, R. Sasidharan.

In Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019.

Measurement of Users' Experience on Online Platforms from their Behavior Logs

[\[Paper\]](#)

D. Jain, A. Sinha, N. Sheoran, D. Gupta, S. Khosla.

In Advances in Knowledge Discovery and Data Mining, PAKDD 2018.

GRANTED PATENTS (5)

Query-oriented approximate query processing based on machine learning techniques.

[\[Patent\]](#)

S. Mitra, N. Sheoran, A. Rao, T. Mai, S. Nair, S. Vaithyanathan, T. Jacobs, G. Siddharth, J. Varshney, V. Maddukuri, L. Mishra.

Machine-learning models applied to interaction data for determining interaction goals and facilitating experience based modifications to interface elements in online environments.

[\[Patent\]](#)

N. Sheoran, N. Raju, V. Srivastava, N. Golakiya, D. Singal, D. Jain, A. Sinha.

Characterizing and Modifying User Experience of Computing Environments Based on Behavior Logs.

[\[Patent\]](#)

D. Jain, A. Sinha, D. Gupta, N. Sheoran, S. Khosla, R. Sasidharan.

Machine-learning models applied to interaction data for facilitating experience based modifications to interface elements in online environments.

[\[Patent\]](#)

A. Sinha, D. Jain, N. Sheoran, D. Gupta, S. Khosla.

Predictive analysis of target behaviors utilizing RNN-based user embeddings.

[\[Patent\]](#)

S. Kim, D. Jain, D. Gupta, E. Koh, B. Kveton, N. Sheoran, A. Sinha, H. Bui, C. Chen

PUBLISHED PATENTS (4)

Causal Multi-touch Attribution.

[\[Patent\]](#)

A. Aggarwal, N. Sheoran, G. Sinha.

Workload Equivalence Class Identification For Resource Usage Prediction.

[\[Patent\]](#)

N. Sheoran, S. Mitra.

Self-learning scheduler for application orchestration on shared compute cluster.

[\[Patent\]](#)

S. Mitra, N. Sheoran, S. Subha, N. Dhake, R. Nehra, R. Simha.

¹Equal Contribution