

Yamuna Krishnamurthy

http://yamunak.com
contact@yamunak.com

- INTERESTS** Developing and applying deep learning, machine learning, data mining, and optimization techniques to inference on graphical, textual, transactional, behavioral and sequence data.
- EDUCATION**
- Royal Holloway University of London, Egham, Surrey, England** **Sep 2018-Present**
PhD Candidate in Computer Science-Machine Learning
Advisor: **Professor Chris Watkins**
- Carnegie Mellon University, Pittsburgh, PA, USA** **Aug 2010-May 2012**
MS in Mechanical Engineering-Machine Learning and Robotics
Advisor: **Professor Jeff Schneider**
- Washington University in St. Louis, MO, USA** **Aug 1998-Aug 2000**
MS in Computer Science-Distributed Computing
Advisor: **Professor Douglas C. Schmidt**
- National Institute of Technology, Nagpur, India** **Aug 1993-May 1997**
BEng in Computer Science
- EXPERIENCE**
- Royal Holloway University of London, Egham, Surrey, England** **Sep 2018-Present**
Post Graduate Researcher, Center for Reliable Machine Learning
- **Gated modular neural networks for transfer learning and interpretability**
 - Exploring modular deep neural architectures, such as learning with experts, for transfer learning and interpretability.
 - **AI in nurse regulatory decision making**
 - Developing a NLP tool that supports nurse regulatory decision making in complaints about nurses in the US, UK and Australia that helps the case managers explore and assess the complaints more efficiently. The tool processes the complaint and determines its risk level, identifies other similar past cases that provide case managers with precedence and additionally matches the complaints with the standards that it is not compliant with. [ACL-Demo 2021]
 - **Computational Immunology**
 - Explored computational models and built an agent based simulation of the T-cell response to infections in the human immune system. Learnt and validated parameters of the simulation model, both hidden and observable, with parameters learnt from experimental data.
- University of Nottingham, Nottingham, England** **June 2019-July 2019**
Data Science Consultant
- **Understanding how policy makers interact with the truth**
 - Goal of the project was to identify and extract references to sources of information by members of the United Kingdom parliament (MPs), from the transcripts of the Hansard policy debates in order to assess the MPs awareness of reliable sources. Fine tuned Google's BERT model with known sources such as reports from the Hansard debates. Used the trained model to then successfully identify other sources in the Hansard debates such as letters, Acts, memorandums and Bills.
- NYU Tandon School of Engineering, Brooklyn, NY** **Feb 2015-Aug 2018**
Research Associate, Visualization Imaging and Data Analysis (VIDA) Center
- **DARPA D3M project**
 - Developing sequence model based automated data mining pipeline generation through self play, using deep neural networks and Monte Carlo Tree Search (MCTS) with reinforcement learning. [ICML-AutoML 2018][ICML-AutoML 2019]
 - **DARPA Memex project**

- Lead researcher and developer of Domain Discovery Tool (DDT), an interactive search and domain specific web content retrieval tool that enables exploratory data analysis of text data, for subject matter experts. [KDD-IDEA 2016]
- Besides technical challenges, the project posed additional challenges that involved communicating with and building an easy to use system for non-data scientists.

Microsoft Research (MSR) and Cloud AI Group, Redmond, WA **Jun 2014-Oct 2014**
Research Intern. Mentors: Dr. Misha Bilenko, Dr. Ofer Dekel, Dr. Yael Dekel, Dr. Rich Caruana

- **Collaborative hyperparameter tuning**
 - Automate hyperparameter tuning for various learners provided by Azure ML (Machine Learning) cloud computing infrastructure.

Technische Universität Dortmund, Germany **Jul 2012-Jun 2014**
Research Associate, Artificial Intelligence Group

- **ViSTA-TV - Video Stream Analytics for viewers in the IP-TV industry**
 - Provide personalized recommendations of TV programs to IP-TV viewers of BBC and Zattoo using a hybrid model and memory based approach of collaborative filtering.

Robotics Institute, Carnegie Mellon University, PA **Apr 2011-Jun 2012**
Research Assistant, Auton Lab

- **Bayesian optimal active search**
 - Detect as many members of a given class, deemed important, as quickly as possible by defining a natural utility function and deriving the Bayesian optimal policy for the associated decision problem. [KDD-MLG 2011] [ICML 2012]
- **Anomaly detection and predictive modeling for counter-terrorism**
 - Developed a predictive model to predict the probability of a socio-political event in the near future based on events in recent past around the world and analyzed the effectiveness of prediction based on lift scores calculated using a sliding time window cross-stream analysis.

University of Bonn, Bonn, Germany **May 2011-Aug 2011**
Visiting Researcher

- **Online pattern sampling and predictive modeling from data streams**
 - Sampled ‘interesting’ patterns from data streams in an online fashion, built an online predictive model and applied on discrete data streams for the uniform and frequency-based distributions.

Goldman Sachs, New York, NY **Apr 2006-Jul 2010**
Information Technology Consultant

- **Repo and Fixed Income Currencies and Commodities (FICC) teams**
 - Designed and developed high performance trade booking and processing applications for quantitative strategists, traders, and sales.

Carlin Equities Group, NJ **Feb 2004-Apr 2006**
Software Consultant

- **Efficient market data distribution and equities trade booking**
 - Designed and developed a framework to subscribe to market data, execute a trading strategy/model against the data and send orders to exchanges using the Financial Information Exchange (FIX) protocol.

OOMWorks LLC, NJ **Sep 2000-Feb 2004**
Co-Founder

- **Distributed Dynamic Scheduling Framework. Client: DARPA**
 - Designed and implemented the Dynamic Scheduling Framework based on Real-Time CORBA 2.0 Specification to enable developing distributed real-time and embedded systems whose Quality of Service (QoS) can be assured even in the face of changes in available resources. [RTAS 2004] [JCBS 2004]

- **Unmanned Air Vehicle (UAV) Video Dissemination.** *Client: BBN Technologies*
 - Used BBN Technologies' Quality Objects (QuO) and Quality of Service (QoS) enabled TAO Audio/Video (A/V) Service for UAV video distribution to achieve predictable and adaptive end-to-end QoS under constrained network bandwidth. [OM 2001] [DOA 2001] [Middleware 2003]

Washington University in St. Louis, MO

Aug 1998-Aug 2000

Student Research Assistant, Distributed Object Computing Lab

- **QoS in TAO's A/V Streaming Service**
 - Integrated ACE QoS framework into TAO's A/V Streaming Service to enable streaming prioritized video and audio data using network QoS reservations through IntServ and Diff-Serv models to provide end-to-end QoS. [IC 2000][DOE 2000]

Motorola Research Lab, Schaumburg, IL

May 1999-Aug 1999

Research Intern. Mentor: Ronald G. Akers

- **Wireless CORBA**
 - Evaluated CORBA mechanisms and services in a wireless environment using mobile IP.

TECHNICAL SKILLS

Programming Languages: Python, Java, C++, Julia, C#, C, JavaScript.

Data Management: MS SQL, MySQL, Cassandra, Elasticsearch.

Data Analytics: scikit-learn, Pandas, StatsModel, Azure ML, Weka.

Deep Learning: PyTorch, Keras, TensorFlow.

Data Visualization: D3.js, React.

Cluster Computing: Microsoft HPC, HPC.

Cloud Computing: Amazon Web Services (AWS).

Tools: Matlab, ACE (Adaptive Communication Environment), TAO (The ACE Orb).

Methodologies: Design Patterns, Object Oriented Programming, Multi-threaded Programming.

INVITED TALK

Flexible and Adaptive QoS Control for Distributed Real-time and Embedded Applications, *Center for Collaborative Control of Unmanned Air Vehicles, University of California, Berkeley, Feb 2009.*

ACTIVITIES

- **Teaching Assistant** for “Deep Learning Course” CS5950 at RHUL, **Spring 2019, 2020, 2021.**
- **Teaching Assistant** for “NLP Course” CS4990 at RHUL, **Spring 2021.**
- **Teaching Assistant** for “AI Principles and Techniques Course” CS5960 at RHUL, **Fall 2019.**
- **Student Representative** on Computer Science Post Graduate Research Student Staff Committee **2019, 2020, 2021.**
- Co-chair, Workshop on Parallel, Distributed and Federated Learning, **ECML PKDD 2020, 2021.**
- Co-chair, Workshop on Decentralized Machine Learning at the Edge, **ECML PKDD 2018, 2019.**
- Program Committee Member, **ECML PKDD 2019.**
- Proceedings Chair, **ECML PKDD 2013.**

HONORS AND AWARDS

- Among 300 out of 1200+ applicants selected (and attended) for the Deep Learning and Reinforcement Learning Summer School 2019, Edmonton, Canada. <https://dlrlsummerschool.ca/>
- 3rd Prize (among 70 teams) at NYU Tandon School of Engineering Research Expo 2017. <https://technical.ly/brooklyn/2017/04/25/nyu-tdandon-research-expo/>
- Mentored the 2nd Place winning team in End Human Trafficking Hackathon, 2016, organized by Manhattan District Attorney's (DANY) office in partnership with Cornell Tech. <http://engineering.nyu.edu/news/2016/11/03/student-hackers-team-manhattan-da-fight-human-trafficking>
- Among top 5 in 1st year undergrad board exams. Passed with distinction and was awarded cash prize.

LANGUAGES

English, Hindi, Telugu, Tamil, Kannada, beginner German.

REFERENCES

Available upon request.

PUBLICATIONS

- Piyawat Lertvittayakumjorn, Ivan Petej, Yang Gao, **Yamuna Krishnamurthy**, Anna Van Der Gaag, Robert Jago, and Kostas Stathis. **Supporting Complaints Investigation for Nursing and Mid-wifery Regulatory Agencies**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 81–91, Online, August 2021. Association for Computational Linguistics (ACL-Demo 2021)
- Iddo Drori, **Yamuna Krishnamurthy**, Raoni de Paula Lourenco, Remi Rampin, Kyunghyun Cho, Claudio Silva, and Juliana Freire. **Automatic Machine Learning by Pipeline Synthesis using Model-Based Reinforcement Learning and a Grammar**. In *AutoML Workshop at ICML*, 2019 (ICML-AutoML 2019)
- Iddo Drori, **Yamuna Krishnamurthy**, Remi Rampin, Raoni de Paula Lourenco, Jorge Piazzentin Ono, Kyunghyun Cho, Claudio Silva, and Juliana Freire. **AlphaD3M: Machine Learning Pipeline Synthesis**. In *AutoML Workshop at ICML*, 2018 (ICML-AutoML 2018)
- **Yamuna Krishnamurthy**, Kien Pham, Aecio Santos, and Juliana Freire. **Interactive Exploration for Domain Discovery on the Web**. In *IDEA Workshop at KDD*, 2016 (KDD-IDEA 2016)
- Roman Garnett, **Yamuna Krishnamurthy**, Xuehan Xiong, Jeff Schneider, and Richard Mann. **Bayesian optimal active search and surveying**. In *Proceedings of the 29th International Conference on Machine Learning*, 2012 (ICML 2012)
- Roman Garnett, **Yamuna Krishnamurthy**, Donghan Wang, Jeff Schneider, and Richard Mann. **Bayesian optimal active search on graphs**. In *Ninth Workshop on Mining and Learning with Graphs, Knowledge Discovery and Data Mining Conference*, 2011 (KDD-MLG 2011)
- **Yamuna Krishnamurthy**, Irfan Pyarali, Christopher Gill, Louis Mgeta, Yuanfang Zhang, S Torn, and Douglas C Schmidt. **The Design and Implementation of Real-Time CORBA 2.0: Dynamic scheduling in TAO**. In *Proceedings of 2004 IEEE Real-Time and Embedded Technology and Applications Symposium*, 2004 (RTAS 2004)
- Christopher Gill, Louis Mgeta, Yuanfang Zhang, Stephen Torri, **Yamuna Krishnamurthy**, Irfan Pyarali, and Douglas C Schmidt. **Enhancing Adaptivity via Standard Dynamic Scheduling Middleware**. *Journal of the Brazilian Computer Society*, 10(1):19–30, 2004 (JCBS 2004)
- Richard E Schantz, Joseph P Loyall, Craig Rodrigues, Douglas C Schmidt, **Yamuna Krishnamurthy**, and Irfan Pyarali. **Flexible and adaptive QoS control for distributed real-time and embedded middleware**. In *Proceedings of the 2003 ACM/IFIP/USENIX international Conference on Middleware for Distributed Systems Platforms*, 2003 (Middleware 2003)
- **Yamuna Krishnamurthy**, Vishal Kachroo, David A Karr, Craig Rodrigues, and Douglas C Schmidt. **Integration of QoS-enabled distributed object computing middleware for developing next-generation distributed application**. In *Proceedings of 2001 ACM SIGPLAN Workshop on Optimization of Middleware and Distributed Systems*, 2001 (OM 2001)
- David Karr, Craig Rodrigues, **Yamuna Krishnamurthy**, Irfan Pyarali, and Douglas C Schmidt. **Application of the QuO quality-of-service framework to a distributed video application**. In *Proceedings of the 3rd International Symposium on Distributed Objects and Applications*, 2001 (DOA 2001)
- Sumedh Mungee, Nagarajan Surendran, **Yamuna Krishnamurthy**, and Douglas C Schmidt. **The Design and Performance of a CORBA Audio/Video Streaming Service**. *Multimedia Networking: Technology, Management, and Applications* edited by Mahbubur Syed, Idea Group Publishing, Hershey, USA, 2001 (MN 2001)
- Douglas C Schmidt, Vishal Kachroo, **Yamuna Krishnamurthy**, and Fred Kuhns. **Developing next-generation distributed applications with QoS enabled DPE middleware**. *IEEE Communications Magazine*, edited by Abdi Modarressi and Sheshadri Mohan, 38(10):112–123, 2000 (IC 2000)
- Vishal Kachroo, **Yamuna Krishnamurthy**, Fred Kuhns, Ronald G Akers, Pradeep Avasthi, Surender Kumar, and Vidya Narayanan. **Design and Implementation of QoS Enabled OO Middleware**. In *Proceedings of the Joint Internet2/DOE QoS Workshop*, 2000 (DOE 2000)