

EDUCATION

Georgia Institute of Technology

Ph.D. in Machine Learning

Atlanta, GA, USA

Aug '21 – Present

University of Toronto

Master of Science in Applied Computing, Dept. of Computer Science; GPA: 3.88/4.0

Toronto, Canada

Sept '18–Dec '19

Indian Institute of Technology (IIT) Guwahati

Bachelor of Technology in Mathematics and Computing; GPA: 8.97/10.0

Guwahati, India

Jul '14–Jun '18

KEY COURSES

- **U of Toronto:** Deep Reinforcement Learning (**A+**), Statistical Learning Theory (**A+**), Machine Learning and Data Mining (**A+**), Spoken Language Processing (**A+**), Machine Learning for Health (**A**)
- **IIT Guwahati:** Data Structures & Algorithms, Theory of Computation, Databases, Data Mining, Probability Theory and Random Processes, Optimization, Scientific Computing, Multivariable Calculus, Matrix Computations, Statistical Inference, Advanced Statistical Algorithms, Real Analysis

EXPERIENCE

Kindred AI

Machine Learning Engineer

Toronto, Canada

Aug 2020 - Aug 2021

- Worked on **grasp-candidate generation** for piece-picking **robotic arms**, designed for e-commerce order fulfillment. Improved the grasping ability of robotic arms using **fully-convolutional grasping** as well as state-of-the-art image segmentation.

Vector Institute

Student Researcher with Prof. Jimmy Ba (University of Toronto)

Toronto, Canada

Apr 2019 - July 2021

- Created **Clockwork VAE**, a **hierarchical latent dynamics model** with temporal abstraction, trained using a variational objective, for long-horizon **video prediction**. Proposed model has a carefully designed filtering posterior representation from observations which gives the model an intrinsic motivation to push shared information in sequences up in the hierarchy. Paper accepted at **NeurIPS 2021**.

Surgical Safety Technologies

Deep Learning Research Intern with Prof. Frank Rudzicz (University of Toronto)

Toronto, Canada

May 2019 - Dec 2019

- Worked on **multi-modal representation learning** (i.e. jointly training vision and language models) for captioning surgical videos. Modeled sequences using a variational recurrent neural network.
- Contributed to the Operating Room Blackbox project, as part of the NLP team, which aims to meet the goal of improving patient safety and outcomes by identifying where errors occur in an operating room and suggesting measures on preventing them.

Microsoft

Software Engineering Intern

IDC, Hyderabad, India

May 2017 - July 2017

- Automated hosting of Jenkins built binaries on **Microsoft Azure** database clusters and made them accessible from **Visual Studio Team Services**, Microsoft Azure's enterprise-grade **CI/CD** solution, through extensive server-side feature development in C#, TypeScript, and Java. My work allowed seamless onboarding of numerous cloud customers already using the Jenkins build system onto the VSTS ecosystem.
- Received a **Full-Time Employment Offer** from Microsoft based on the work done.

PUBLICATIONS

- **V. Saxena**, J. Ba, D. Hafner, "Clockwork Variational Autoencoders," **NeurIPS 2021**
- **V. Saxena**, S. Sivanandan, P. Mathur, "Dyna-AIL: Adversarial Imitation Learning by Planning," *Beyond "Tabula Rasa" in Reinforcement Learning Workshop*, **ICLR 2020**
- **V. Saxena**, N. Yadala, R. Chourasia, F. C.-H. Rhee, "Type Reduction Techniques for Two-dimensional Interval Type-2 Fuzzy Sets," *IEEE International Conference on Fuzzy Systems (FUZZ-IEEE)*, 2017
- R. Chourasia, **V. Saxena**, N. Yadala, F. C.-H. Rhee, "Visualization of Two-dimensional Interval Type-2 Fuzzy Membership Functions using General Type-2 Fuzzy Membership Functions," *Joint 17th World Congress of International Fuzzy Systems Association and 9th International Conference on Soft Computing and Intelligent Systems (IFSA-SCIS)*, 2017

PROJECTS

Video Prediction with Variational Temporal Hierarchies

[arXiv][Github]

Research project with Prof. Jimmy Ba (U of Toronto) and Danijar Hafner (Google Brain, Toronto)

May 2019 - Feb 2020

- Created **Clockwork Variational Autoencoder**, a temporally abstract latent dynamics model for **video prediction** which makes use of **temporal hierarchies to predict for over 1000 frames into the future**. The model consists of a filtering posterior which encodes different amounts of temporal information at each level, and a prior which can predict into the future without ever being directly conditioned on the observation frames, and is trained using the **ELBO objective**.
- Also showed separation of information at different levels of the hierarchy: with moving MNIST, the digit identities were stored at higher levels of the hierarchy than digit movements; with the GQN mazes dataset, the color and texture of the walls were stored at higher levels than camera movements.

Dyna-AIL: Adversarial Imitation Learning by Planning

[arXiv]

Research project with Prof. Jimmy Ba (U of Toronto)

Sept 2018 - Dec 2018

- Proposed an **end-to-end differentiable adversarial imitation learning algorithm in a Dyna-like framework** for switching between model-based planning and model-free learning from expert data. We computed stochastic value gradients of the discriminator's output w.r.t. states and actions over entire trajectories, used the reparameterization trick to propagate gradients through the dynamics model, and optimized the policy network using TRPO.
- Our results on both discrete and continuous MuJoCo simulation environments showed that our approach **converges to an optimal policy with fewer number of environment interactions** in comparison to the state-of-the-art model-based reinforcement learning methods such as MAIL [Baram et al., 2017].
- Implemented in Python using TensorFlow.

Evaluating GANs using Real-World Estimate

[Report] [Github]

Research project with Prof. Roger Grosse (U of Toronto)

Sept 2018 - Dec 2018

- Proposed a MCMC sampling method using the Metropolis-Hastings algorithm for obtaining **samples from the 'real distribution' as learned by the discriminator of a GAN**, by using the generator network as a proposal distribution. Improved upon an existing rejection sampling approach proposed by Azadi et al. (2018) with respect to a **higher acceptance percentage and less samples rejected around the modes**. Also showed how our method can be used to evaluate a metric containing a reverse-KL divergence component.
- Implemented in Python using TensorFlow.

fastMRI: Enhancing MRI reconstruction through robust k-space interpolation

[Report]

Research project with Prof. Marzyeh Ghassemi (U of Toronto)

Jan 2019 - Apr 2019

- Worked on MRI reconstruction using undersampled k-space to generate accurate high-quality images. Proposed a **k-space imputation technique using de-noising autoencoders**, which outperformed the majorly accepted zero-filling reconstruction baseline.
- Implemented in Python using PyTorch.

Time Series Modeling of Bitcoin Prices

B.Tech. Thesis Project with Prof. Arabin Kr. Dey (Dept. of Mathematics, IIT Guwahati)

July 2017 - Apr 2018

- Built a time series model that predicts bitcoin prices to assist in cryptocurrency trading. Explored various auto-regressive and moving average statistical time series models, and also others for modelling volatility such as ARCH and GARCH models. Worked on an **ensemble of LSTM Encoder-Decoder and statistical time series models** to predict univariate and multivariate time series.
- Also used MCMC based **Approximate Bayesian Computation** to estimate parameters of non-linear univariate and bivariate time series models, and was able to obtain state-of-the-art signal accuracy for cryptocurrency prices.
- Implemented in Python using TensorFlow.

ACHIEVEMENTS & AWARDS

- Addictive Mobility Scholarship in Applied Computing:** Awarded by University of Toronto to top 4 students based on academic merit.
- KVPY Fellow Award:** Awarded the Kishor Vaigyanik Protsahan Yojana 2013 Fellow Award with an All India Rank of 178.
- Indian National Mathematics Olympiad:** Qualified the Regional Mathematics Olympiad '13 with State Rank 25 and appeared for INMO.
- JEE Advanced:** Secured All India Rank 1563 (99.9 percentile) in Joint Entrance Examination 2014.
- Microsoft Code.Fun.Do '16:** Won the First Runner-Up prize for building an AI powered assistant for Microsoft Word Online.

TECHNICAL SKILLS

- Languages:** Python, C/C++, C#, Java, TypeScript
- Numerical Computation:** TensorFlow, PyTorch, Matlab
- Databases:** MySQL, MongoDB, Neo4j