

Valliappa Chockalingam

E-mail: valli@umich.edu

Personal Website with source code, papers, and more information: <http://geekstor.github.io>

Education

University of Alberta

Edmonton, AB

Master of Science in Computing Science

Sep. 2018 – Present

Supervised by: Professor Richard Sutton

- Courses: Machine Learning, Artificial General Intelligence, Reinforcement Learning, Stochastic Approximation

University of Michigan

Ann Arbor, MI

Bachelor of Science in Engineering in Computer Science

Sep. 2013 – May 2017

- Notable Courses: Machine Learning, Artificial Intelligence, Computational Neuroscience, Agent Based Modeling, Game Theory, Theory of Complex Systems, Cryptography, Computer and Network Security, Computer Architecture and Organization, Theory of Computation, Programming and Data Structures, Data Structures and Algorithms
- Minor in Complex Systems

Publications

Control of Memory, Active Perception, and Action in Minecraft

ICML '16

J. Oh, V. Chockalingam, S. Baveja, H. Lee

Extending World Models for Multi-Agent RL in MALMÖ

AIIDE '18

V. Chockalingam, T. Sung, F. Behbahani, R. Gargeya, A. Sivanantham, A. Malysheva

Experience

Deep Learning Camp Jeju

Jeju, South Korea

Participant

Jun. 2018 – Aug. 2018

Acceptance Rate: 4.2% (24 / 562 applicants from 50+ countries)

Working with: Rishab Gargeya, Yu-Han Liu, and Taehoon Kim

- Implementing Distributional Reinforcement Learning techniques
- Extending on existing techniques by adding risk-adaptivity

Vee Technologies

New York, NY

Technology Operation Specialist

Oct. 2017 – May. 2018

- Worked on developing Business Process diagrams to understand chart retrieval and off-shore coding processes
- Helped with project to automate ICD code extraction from charts to help coders work faster

Preferred Networks

Tokyo, Japan

Research Intern

Jul. 2017 – Oct. 2017

Supervised by: Toshiki Kataoka and Brian Vogel

- Read papers on Multiagent and Multitask Deep Reinforcement Learning
- Worked on general game playing with the “rules” for tasks provided in a Domain Specific Language
- Conducted experiments with novel A3C-inspired agent architectures looking at generalization ability

University Of Michigan

Ann Arbor, MI

Directed Independent Study Researcher

Sep. 2016 – Dec. 2016

Supervised by: Professor Satinder Singh Baveja

- Read papers in hierarchical reinforcement learning and planning
- Developed Value Iteration Network inspired agents that can plan at different temporal scales
- Wrote a report discussing results and implications of the constructed hierarchical planner

EECS 445 (Machine Learning) Instructional Aide

Sep. 2016 – Dec. 2016

- Created programming assignments for homework and hands-on lecture sections
- Taught weekly discussion and hands-on lecture sections

Microsoft Research

Cambridge, UK

Research Intern

May 2016 – Aug. 2016

- Helped develop the Minecraft AI Platform, Project Malmö, along with a variety of tasks

- Implemented a range of RL agents (primarily in TensorFlow and Chainer), from traditional RL agents like SARSA- λ to more complex, then state-of-the-art, Deep RL agents like Dueling DQN
- Looked at generalization performance of agents using difficulty metrics based approaches

University Of Michigan

Ann Arbor, MI

EECS 545 (Graduate Machine Learning) Instructional Aide

Jan. 2016 – Apr. 2016

- Created Jupyter notebook based lecture slides using Markdown and LaTeX
- Wrote Python programs to demonstrate various Machine Learning concepts

Machine Learning Research Assistant

Jun. 2015 – Dec. 2016

Working with: Professor Satinder Singh Baveja, Professor Honglak Lee, Junhyuk Oh

- Developed memory based Deep Q-Network agents that can learn to traverse “real-world” like 3-Dimensional environments and perform complex tasks with good generalization capabilities
- Deepened understanding of Deep Learning and Reinforcement Learning research by reading relevant papers
- Gained experience with relevant programming languages such as Java, Python, and Lua, and Deep Learning packages such as Torch

EECS 203 (Discrete Mathematics) Instructional Aide

Jan. 2014 – May 2015

- Managed and answered questions on the online Piazza forum
- Conducted weekly office hours and discussion sections
- Created “challenge” problems for students keen on going “beyond the syllabus”

EdGE Networks

Bangalore, India

Intern

Apr. 2014 – Jun. 2014

- Learned Natural Language Processing techniques
- Worked on developing recommender systems that match résumés to companies through web scraping
- Gained experience with Python packages such as Scikit-Learn, PyBrain, NLTK, and BeautifulSoup

iVista Solutions

Bangalore, India

Intern

Jun. 2012 – Aug. 2012

- Learned to create and manage simple web servers
- Worked with PHP, MySQL, HTML, CSS, and Wordpress as a CMS
- Used what was learned to create and manage a website for a cancer fundraiser marathon

Taneja Aerospace

Hosur, India

Intern

Apr. 2012 – Jun. 2012

- Constructed models for remote controlled airplanes using CAD software
- Used wind tunnels and computer simulations to test the aerodynamics of different models
- Built and tested a RC plane model on an actual runway

Awards

University of Michigan

Ann Arbor, MI

Dean's List

Fall 2013, 2014, 2015, 2016, Winter 2017

University Honors

Fall 2013, 2014, Winter 2017

Indus International School

Bangalore, India

Young Achiever Award for Academic Achievement

Winter 2013

Gear Innovative International

Bangalore, India

Mini-Nobel Prize Nominee

Summer 2011

Languages

Programming and Markup Languages, and Frameworks: Python, Chainer, PyTorch, TensorFlow, Keras, NumPy, SciPy, Matplotlib, Java, C++, C, LaTeX, Markdown, NetLogo, Lua, Torch, MATLAB, PHP, MySQL, HTML, XML, CSS, JavaScript, Swift, Objective-C, R, Verilog

Languages: English, Tamil, Japanese, Spanish, French, Kannada

Projects

Evaluating Generalization with Distributional Reinforcement Learning*: Developing a risk-adaptive Deep Reinforcement Learning agent that can explore better using knowledge of its uncertainty in its estimated value distribution for better and faster transfer to new tasks and domains.

General video game playing with descriptions*: Using multi-modal Deep RL architectures to facilitate feature extraction from states, together with game descriptions, using an asynchronous advantage actor-critic training scheme for fast multitask learning with language as an inductive bias.

Reinforcement learning agents for Riichi Mahjong*: Assuming a POMDP formulation, tracking belief states and the fact that opponents cannot win with tiles they have discarded to create RL agents that can learn to play Riichi Mahjong to human or superhuman levels despite high partial observability.

Hierarchical differentiable planning networks: Extended on Value Iteration Networks with options and dilated convolutions, and analyzed inferred reward and value functions.

Studying depression and reward processes through Serotonergic neurons: Created computer models of Serotonergic neurons in the Dorsal Raphe Nucleus and studied spiking behavior to understand how various parameters could affect reward processing and cause or cure depression.

Anomaly detection in Controller Area Network messages: Looked at how anomalies in the messages traversing vehicles through CAN buses can be detected through stateless one-class Support Vector Machines and stateful time-series based Long Short Term Memory neural networks.

Deep Q-Network based agents with external memory in Minecraft: Used Minecraft Forge, Lua and Torch, to develop novel memory based Deep Q-Network agents that can better learn to solve and generalize over different challenging tasks like goal-search in 3-Dimensional partially observable mazes compared to prior baseline agents.

Incentivizing exploration for learning to play games with randomness: Built Deep Q-Network agents where exploration is incentivized using the error between true and predicted next state encodings as a reward bonus, with the current state and ϵ -greedily chosen action as input to a denoising autoencoder, allowing for better handling of random aspects of domains, like enemies in Atari 2600's Seaquest.

Detecting impurities through ultrasonic interferometry: Looked at new ways of testing for adulteration in foods and liquids using ultrasonic interferometers to help identify harmful (and sometimes carcinogenic) impurities like Calcium Carbide in Mangoes and Urea in Milk

Search engine keyword popularity prediction: Used NetLogo to simulate system dynamics of a large network along with real world statistics obtained from Google Analytics to try to replicate patterns in keyword popularity.

Social network structure analysis: Extended on Schelling's model of segregation with new computer models to explain how social groups are formed and how the preferences of individuals affect underlying social network structure.

Sociopolitical opinion analysis: Used different complex systems models such as Voronoi maps to find social and political scenarios where the opinions of a group do better than the opinion of a few experts.

* In Progress