

Srinivas Venkattaramanujam

+91 77082 09071 • Email: srinivas.vr.edu@gmail.com

srinivr.github.io • linkedin.com/in/srinivas-venkattaramanujam-95560a62

EDUCATION

MILA & McGill University

PhD in Computer Science

Montreal, Canada

Aug 2020 - Aug 2021 (Dropped out)

MILA & McGill University

Master of Computer Science - Thesis

Montreal, Canada

Sep 2017 - Aug 2020

Thiagarajar College of Engineering

B.Tech in Information Technology

Madurai, India

Aug 2010 - May 2014

SKILLS

Areas

Deep Learning, Reinforcement Learning, Automatic Speech Recognition, Backend development (Java/Servlets/JSP), Android development, Mathematics - Real Analysis, Linear Algebra, Probability and some Topology

Languages & Frameworks

PyTorch, Kaldi, OpenFST, SRILM, Java, Python, C, MySQL, Redis, Elasticsearch, Android, Git

SOFTWARE ENGINEERING EXPERIENCE

Phaidra.AI

Software Engineer

Remote

Aug 2021 - Present

- Built and maintained microservices, contributed to ML codebase
- Identified and resolved bottlenecks in data processing
- Worked on improving MLOps stack

Pickyourtrail.com

Senior Development Engineer

Chennai, India

Jul 2015 - Feb 2016

- Developed several APIs for a personalized itinerary planning tool
- Developed an Android app to display itineraries for individual customers

ShopperLane

Owner

Chennai, India

Feb 2015 - Jul 2015

- Developed an Android application to search for products in neighbourhood stores
- Developed a suite of web applications for inventory management, reporting and billing

Verizon Data Services India

Software Engineer

Chennai, India

Aug 2014 - Feb 2015

- Worked on the backend development of a quoting tool

RESEARCH EXPERIENCE

MILA & McGill University

Research Assistant

Montreal, Canada

Jan 2018 - Present

- Working on Reinforcement Learning and Representation Learning with **Prof. Doina Precup**
- Primary research objective is to eliminate the need for domain knowledge by RL agents

Speech Lab, IIT Madras

Project Associate

Chennai, India

Feb 2016 - Apr 2017

- Worked on Automatic Speech Recognition (ASR) with **Prof. Umesh Srinivasan** • Used Deep Learning to build Automatic Speech Recognition (ASR) systems and investigated the use of distillation for cross-lingual transfer in ASR for resource constrained languages

PUBLICATIONS

- **Venkattaramanujam, S., Crawford, E., Doan, T., & Precup, D. (2020, February). Self Supervised Learning Of Distance Functions For Goal Conditioned Reinforcement Learning. *Preprint*. [Paper]**

We propose an approach to learn state embeddings that are useful for goal-conditioned policies. We discuss the conditions that guarantee the existence of this embedding space and propose a practical approach to approximate it.

PROJECTS

- **dqn-pytorch-lib**, an implementation of DQN and n-step DQN using **PyTorch**. Implemented the TreeQN architecture and reproduced the results. The implementation supports modular addition of auxiliary losses such as reward prediction errors, state prediction errors and so on.
- **tf-kaldi**, a binding between Tensorflow and Kaldi for ASR. The DNN component of the acoustic model is implemented in Tensorflow whereas the feature extraction and decoding is handled using Kaldi.
- **kaldi-long-audio-alignment**, a toolkit to split a long audio and the corresponding transcription into multiple non-overlapping shorter segments. The correct transcription for the shorter audio segments are automatically determined from the transcription of the long audio.
- **fast-speech-transcriber**, a tool to enable easy transcription of audio files. The tool automatically loads the audio files in a given directory and provides controls to control audio playback and to navigate the audio files. The transcriptions are saved automatically. The tool uses a dictionary to provide auto-complete suggestions as the words are being typed.
- **A Guided tour of 'Metrics for MDPs with Infinite State Spaces'** (theory), provides the mathematical background of and the explanation of the proofs in *Metrics for MDPs with Infinite State Spaces* - a state aggregation method for Markov Decision Processes. The state aggregation is performed using the extension of the notion of bisimulation to a metric in continuous state spaces, called the bisimulation metric.

GRADUATE COURSES

Machine Learning, Probabilistic Graphical Models, Reinforcement Learning, Matrix Computations, Representation Learning and Mathematical Foundations of Machine Learning