2020 Science Atlantic Mathematics, Statistics, and Computer Science Meeting

Science Atlantic
connecting science education and research

Hosted online by the Department of Mathematics and Computing Science

October 31-November 1 2020

	SATURDAY OCT 31			SUNDAY NOV 1		
8:45	OPENING REMARKS				Hopkins	**MATH B**
9:05	Zassmann	**CS A**	Clow **MATH A**	Taylor-Melanson **CS B**	Smith	9:05
9:25	Rahman		Liu	Jabre	Lee	9:25
9:45	Pandey		Barton	Zhou	Harrigan	9.45
10:05	Kollo		Porter	Leblanc	Zhang	10:05
10:25	Melanson		Macquarrie	Field Lecture James Watmough 10:30-11:30		
10:45	BREAK			Break		
11:00	Sedgewick Lecture Sreejata Chatterjee 11:00-12:00			Kaur	**GRAD A**	Kandre **GRAD B**
11:30				Ramegowda		Chitla 12:00
11:40				Bussey		Thakur 12:20
12:00				Results, closing remarks		12:45
12:20						
12:40						
1:00			CS Contest (5 hrs)			
2:00	Math Contest (3 hrs)					
2:45		Joint meeting				
3:00						
4:00		Math	CS			
5:00			Ends 6:00			
7:30	Blundon Lecture Gordon Macdonald/Andrew Godbout (UPEI) 7:30-8:30					

Invited Speakers

Sreejata Chatterjee (Sedgewick Lecturer) is the co-founder and Head of Product at LeadSift - a marketing technology company. She's responsible for the design, architecture and shipping of LeadSift products, figuring out how to deliver maximum value to the customers, general operations of the business and also picking up snacks for the office. Outside of work, she's a guest lecturer at Saint Mary's University and for Ladies Learning Code, co-hosts a Data Science meetup group and runs a small-scale non-profit company but you're most likely to find her chilling over wine and food after a stroll that she'll insist is a hike!

Andrew Godbout (Blundon Lecturer) is a computer scientist who was trained at Saint Mary's U. (B.Sc.) and U. of Calgary (M.Sc. and Ph.D.). Born and raised in Dartmouth, NS, after stints at IBM and professional speed-skating, Andrew sought a return to the Maritimes and took a position at UPEI, where he has been since 2016. His main research is in computer vision and human computer interaction, but recently has branched into Machine Learning.

Gordon MacDonald (Blundon Lecturer) is a mathematician who was trained at Dalhousie U. (B.Sc.) and U. of Toronto (M.Sc. and Ph.D.). Born and raised in Cape Breton, after stints at Waterloo and Berkeley, Gordon sought a return to the Maritimes and took a position at UPEI, where he has been since 1993. His main research is in linear algebra and functional analysis, but recently has branched into Machine Learning.

James Watmough (Field Lecturer) is an applied mathematician at the University of New Brunswick. He works on mathematical modeling, with an emphasis on population modeling and epidemiology.

Abstracts

A Medley of Results Regarding Acyclic Polynomials of Graphs

Caroline Barton (Dalhousie)
(work with Jason I. Brown)

An acyclic subset of a graph G with vertex set $V(G)$ is a subset of $V(G)$ that contains no cycles. We call the generating function for the number of acyclic subsets of each cardinality i the *acyclic polynomial* of G. We will investigate acyclic polynomials of low and high degree and investigate how the complex roots behave. Are they always in the left half-plane? Can they have arbitrarily large modulus? Can the limits of roots for families of such polynomials coalesce into curves? We will talk about all of these interesting problems and more.

Eigenvalues of the MOTS stability operator for slowly rotating Kerr black holes

Liam Bussey(MUN)
(work with Graham Cox and Hari Kunduri)

We study the eigenvalues of the MOTS stability operator for the Kerr black hole with angular momentum per unit mass $|a| \ll 1$. We prove that each eigenvalue depends analytically on a, and compute its first nonvanishing derivative. Recalling that $a = 0$ corresponds to the Schwarzschild solution, where each eigenvalue has multiplicity $2\ell + 1$, we find that this degeneracy is completely broken for nonzero a. In particular, for $0 < |a| \ll 1$ we obtain a cluster consisting of ℓ distinct complex conjugate pairs and one real eigenvalue. Moreover, for $\ell \geq 1$ the real eigenvalue is a convex function of a, whereas the $\ell = 0$ mode (the principal eigenvalue) is concave, and hence has a maximum at $a = 0$.

Big Data and Local Startups

Sreejata Chatterjee (Leadsift)

Business meets Data. We will explore a couple general startups and a longer list of local startups that are truly utilizing Big Data and making exciting products, and then follow the journey of a specific business (mine!) that is a bit too close for comfort to Cambridge Analytica and how it is completely "legal" and possible.

An Intelligent Approach to Predict Drug Combinations towards More Effective Treatment of Cancer

Eshitha Reddy Chitla (SFX)
(work with Othman Soufan)

Cancer is a chronic disease with estimates that about one in every two Canadians develops cancer. One of the significant challenges in cancer treatment is its ability to return and develop again, even after treatment. Though the recurrence rates widely depend on cancer types, stages of cancer, genetic factors, and treatment delivered, most types are likely to be recurring. One way to treat recurrent cancer is to use combination drugs therapy. This depends on combining two or more drugs that work according to different mechanisms, thus decreasing the likelihood of developing resistant cancer cells. Advancements of computer science applications in biology, shaped by the field of "bioinformatics", have offered the chance to examine solutions to this challenging problem. Here, we will share our efforts to develop machine learning algorithms (a subclass of artificial intelligence) to recommend effective drug combinations. Our dataset is prepared using one database with 5000 number of samples. In this talk, we aim to explain the steps needed to formulate and tackle such challenging problems using AI.

A New Approach to Finding Ideal Play in Poset Games

Alexander Clow (SFX)
(work with Stephen Finbow)

Poset games are a class of combinatorial game that remain unsolved. Examples of poset game include Nim, Chomp, Subset Take-Away, Divisors and Geography. Wilson and Soltys 2011 proved non-constructively that polynomial time algorithms for finding ideal play on any poset game exists, but with the notable exception of Nim, fast algorithms to find specific ideal strategies are unknown. What is more, formidable computational efforts such as that of Zeilberger 2001 have failed to find any global patterns in even special cases such as 3 by n Chomp. This talk presents original results which classify some positions as winning or losing positions and establish the equivalence of ideal strategies on posets that are seemingly unrelated. We believe these results provide a basis for novel analytical methods of finding ideal play in all poset games.

Volume Preserving Neural Networks

Andrew Godbout and Gordon MacDonald (UPEI)

We give an introduction to the mathematics behind the structure and training of neural networks and the problem of vanishing gradients which makes it difficult to train deep neural networks. We then describe a new architecture we have developed for deep neural networks, one where all layers (except the output layer) are (hyper)volume preserving. This control on the volume helps control the gradients and allows volume-preserving deep neural networks to train reliably, quickly and accurately, and with consistent learning across layers. To demonstrate this we apply our volume-preserving neural network model to some standard datasets.

Kendall's Tau for Time Series Data: Developing a New Way to Test for Correlation

Patrick Harrigan (Acadia)
(work with Ying Zhang)

Kendalls Tau is a non-parametric statistic used on bivariate data that measures whether or not these two variables are independent and, if not independent, assesses the type and degree of dependency that exists between them based on the concordance (and discordance) structure of the data. One of the assumptions of Kendalls Tau procedure requires that the paired observations be mutually independent and identically distributed (iid) according to some, known or unknown, continuous distribution. Requiring this assumption causes certain issues when exploring certain datasets with non-iid structures, particularly, bivariate time-series data. The purpose of this research is to examine the behaviour of Kendalls Tau when the iid assumption is *relaxed* to include weakly stationary time-series data. We will discuss different methods to find a reliable variance estimate to construct accurate confidence intervals when dealing with these types of data. Additionally, we will determine which one of these estimates are best given a certain situation looking at both the advantages and disadvantages to all proposed estimates. Finally, once these methods are developed, we will conduct the data analysis of a time-series dataset that involves assessing the migration patterns and levels of certain genera found in the gut microbiome.

Estimating underreporting of COVID-19 cases using fatality data

Margaret Hopkins (Acadia)
(work with Acadia COVID-19 Modelling Group (Duane Currie, Coleman Hooper, Dr. Richard Karsten, Yifan Li, Franklin Mendivil, Holger Teismann))

Abstract: data. Testing and reporting policies vary across regions, and in most regions it is acknowledged that the reported numbers of COVID-19 cases do not reflect the true incidence of the disease. In this talk, a model is introduced that uses death data and estimates of the infection fatality ratio (IFR) to produce estimates for the true number of COVID-19 cases in various regions. This model was developed over the summer of 2020 to analyze the "first wave" of the COVID-19 pandemic in various regions.

Using Machine Learning to Overcome Selection Bias in Metaheuristics

Kosay Jabre (SFX)
(work with Antonio Bolufe-Rohler)

Metaheuristics have enjoyed success as the method of choice for solving many real world optimization problems due to their flexibility and speed. However, metaheuristics suffer from selection bias due to selection operators using just fitness values when comparing sampled solutions against proposed solutions. This results in a limiting of the exploration of the search space, where proposed solutions from a more promising region (with a more optimal extremum) are more likely to be discarded. Our research aims to utilize deep learning to combat this selection bias by modifying the fitness values used for selection. An encoder neural network is used to encode the current state of the optimization problem. A decoder neural network assigns to a proposed solution an auxiliary fitness that estimates the fitness of the local extremum of the solution's attraction basin. Metaheuristics then optimize an objective function which is a linear combination of the original fitness and this generated auxiliary fitness.

Asymptotic Iteration Method Meets Black-Scholes

Jiaying Liu (UPEI)

(work with Nasser Saad and Kai Liu)

Financial derivatives are becoming increasingly important since 2008 financial crisis where an umbrella group for regulators over the world oversaw derivatives markets. Nonetheless, the regulators are having an incomplete picture of financial derivative market after a decade. The Black-Scholes (BS) model is a well-known method for option pricing. However, the BS model requires too many assumptions which do not hold in the real market. In this presentation, we introduce the Asymptotic Iteration Method to derive a closed form solution for problems in financial mathematics, and then extend the normal BS model to a new model.

Pricing of Multivariate Financial Derivatives

Jingfan Zhang (UPEI)
(work with Alexander Alvarez)

The objective of this work is to provide an efficient methodology for the accurate pricing of multivariate financial derivatives. Our methodology uses polynomial approximations obtained using a least squares criteria. We made the theoretical developments under a multivariate Black-Scholes model, and implemented this approach in R for both the bivariate and the trivariate cases. Our numerical results indicate that our approach compares favourably to competing Monte Carlo methodologies.

ChemToxXplorer: A web tool to explain and visualize the predictions of chemical toxicity

Shree Nikhi Kandre (SFX)
(work with Othman Soufan)

Regulatory agencies like Environment Canada and toxicological research communities are interested in developing, testing, and applying new approaches to evaluate chemical hazards more efficiently. Given the complexity of analyzing thousands of chemicals simultaneously, there is a need to develop automated solutions to make the process more efficient. Here we introduce ChemToxXplorer as a unique online platform integrating diverse resources to systematically characterize chemical structures associated with different toxicity endpoints. ChemToxXplorer is a high-performance cloud-based web application with an intuitive interface to allow users to easily browse, search, analyze, and visualize their results. The underlying knowledge base is based on data curated from Tox21, ACToR, CTD, T3DB, LINCS L1000, and Open TG-GATEs with links to PubChem as well as other compound databases. ChemToxXplorer is designed for toxicologists, biologists, biochemists, and clinicians interested in getting useful information about possible adverse effects for a set of provided chemicals. We plan to use Google compute engine to host our platform on the cloud. This talk will highlight our current efforts in building a web system that will contribute to predicting toxic chemicals and support a wide variety of applications in environmental toxicology.

Developing an Artificial Intelligence Solution to Study Disease-Food-Drug Interactions

Jaskiran Kaur (SFX)
(work with Othman Soufan)

Several studies have addressed the effects of food intake on prescribed drugs, especially for patients with chronic diseases. For example, following an inappropriate diet program has been shown to lead to potential negative consequences (a.k.a. adverse effects) in cardiovascular, obesity, and other types of diseases. This is due to specific food ingredients (e.g., artificial sweeteners and colors, emulsifiers) that may impact the process of drug releases, absorption, or metabolism. Here, we explore the application of artificial intelligence (AI) in studying the adverse effects of disease-food-drug interactions. This type of formulation is an example of many applications in the domain of bioinformatics. We will illustrate how a dataset can be constructed with samples representing interactions of disease terms, chemical food ingredients and drug items. Then, we discuss how advanced machine learning models, such as deep learning, can represent a solution to the problem. The dataset we prepare is based on mining DrugBank.ca and FoodDB.ca, two well-known databases developed by The Metabolomics Innovation Centre of Canada. This work will help physicians and pharmacists recommend which food ingredients patients should cautiously take for their drug therapy. The talk will illustrate in general the promising ways to upgrade the access and quality of treatment by outlining AI applications in disease-food-drug datasets.

Dynamic Perspective Constellation Identification UsingPattern Matching and Rotational Invariance

Nikolaus Kollo (SMU)
(work with Jiju Poovvancheri)

Human beings have long used the stars as a mechanism to navigate great distances with an astonishing accuracy. Over generations mankind has mastered this ability and has since moved on to greater scientific endeavours like astronomy, constructing satellites, and space flight. These fields have have thrived utilizing cutting edge technologies like high precision gyroscopes and sensors. Perhaps the stars can once again help us navigate great distances and eliminate the need for precision sensors in these fields. The first steps towards this would be to use computer vision techniques to build a star tracker that can recognize constellations in the night sky. By utilizing pattern matching on high luminosity stars, software could be trained to spot patterns in the sky - like a human would.The first algorithm would be responsible for determining the brightest star visible in the sky, where the subsequent algorithms would then measure its relative distances to other bright stars, and use these distances to create unique

shapes. These shapes would exploit their uniqueness regardless of perspective by rotational invariance and thus could be used to identify star patterns. Once a pattern is confirmed to be a statistical match by one final algorithm, the software can then confidently identify any other visible star in the sky relative to a known point.

Cops, Robbers and Donut Shops: The Implications of Constrained Cop Mobility

Taylor Lee (UNBSJ)
(work with Andrea Burgess)

The premise of pursuit and evasion – effectuated upon geometric formations and graphical structures – possesses significant implications in the context of graph theory. In essence, one may acquire an enriched understanding of various graph classes – and the properties thereof – by considering the manner in which two hypothetical entities interact with a given structure. Accordingly, the game of *Cops and Robbers*, which comprises a set of cops (controlled by one player) who endeavour to capture a robber (controlled by the opposing player), has been the topic of a robust body of work over the preceding forty years. The subsequent presentation shall consider a variation to the traditional game of Cops and Robbers, wherein cop mobility will be constrained by the institution of entrapments – referenced henceforth as *donut shops*. In particular, the ramifications of this variant will be examined for three classes of graphs: paths, cycles, and finite grids. A series of formulae will be presented, in an endeavour to delineate the *cop number* – that is, the minimum quantity of cops necessitated to capture the robber within a finite number of moves – for each of the three respective classes. Consequently, certain general remarks and conclusions will be drawn, in an effort to demonstrate the implications of constrained cop mobility.

Asymptotic Iteration Method for Second Order Linear Hahn Difference Equations

Lucas Macquarrie (UPEI)
(work with Nasser Saad and Shafiqul Islam)

The Hahn difference operator $D_{q;w}f(x) = (f(q\,x + w) - f(x))/((q-1)x + w), q, w > 0$ is a generalization of both the q -derivative and the forward difference operator. We extend and unify the asymptotic iteration method for q -difference equations and discrete difference equations to this operator, deriving a solution for a special case of second order linear homogeneous Hahn difference equations.

Adding the Pseudo-maximality Capability to the Formal Language Server

Patrick Melanson (SMU)
(work with Stavros Konstantinidis)

The Language Server (LaSer) is a website used to ask different questions concerning regular languages. The three primary functions of the website are to decide satisfiability, maximality, as well as constructing languages given certain conditions. When deciding satisfiability, we provide the website a regular expression, as well as either a pre-determined transducer, or one that the user has already created. This transducer represents the quality we wish to satisfy. If the language satisfies the condition, we return true. If not, we return false as well as a counter example, or a witness. If a language satisfies a given property, we can then ask if there are any other words that we can add to the language, that also satisfy the property. If it is not possible to add new words, we say the language is maximal with respect to the given property. However, deciding maximality can become quite expensive in terms of time, so we wish to find a better method. Because the problem of maximality is PSPACE-complete, we believe that there is no fundamentally better algorithm. So instead, if the language or property is too complex, we wish to see if the language is p%-maximal. That is, we pick words at random (via some probability distribution) that are not in our original language, and we check to see if they satisfy our property. If the word does satisfy the condition, we have found a witness, and thus return false. If the word doesn't satisfy the property, we try again. This step repeats until we've either found a witness or we've checked enough words to be 99.99% sure that it is p%-maximal. We call this pseudo-maximality. Checking pseudo-maximality is much more efficient then checking complete maximality, and has the added bonus of not giving a false witness.

Towards video based collective motion analysis through shape tracking and matching

Bivash Pandey (SMU)
(work with Jiju Poovvancheri, Sumesh Thakur, Hemanchal Joshi, Ayusha Pradhanga, and Dr.Yasushi Akiyama)

The term collective is used to refer to a group of animals such as a flock of birds or a herd of elephants. Aggregate motions of such collectives often give rise to visually pleasing shapes and patterns (e.g. V-shape formation of geese while they migrate from one place to another). While shapes of moving collectives are of great interest in many scientific studies, scant attention has been given to algorithmically extract and render these shapes via polygonal boundaries or graphs. Here, we present a multi-stage, proof of concept framework for

tracking geometric shapes and extracting video frames containing a user defined shape of moving collectives, by employing a deep-learning based object detection, well-known alpha shapes and a modified shape context. We demonstrate the usefulness of the proposed framework on a couple of test videos and discuss its potential applications in a wider area.

Hyperopic Cops and Robber

Amanda Porter (MTA)
(work with N. Clarke (Acadia), S. Finbow (StFX), and M. Messinger (MtA))

We explore a variant of the game of Cops and Robber introduced by Bonato et al. where the robber is invisible unless outside the common neighbourhood set of the cops. The hyperopic cop number is the compliment to the cop number and we investigate bounds of this quantity. We define a *small common neighbourhood set* and relate the minimum cardinally of this graph parameter to Hyperopic Cops and Robber. We analyze various diameter 2 graphs, with a particular focus on the join of two graphs. We investigate a general upper bound for the Cartesian product of graphs.

The Mobile App Implementation of Dewblock Clients

Minhajur Rahman (UPEI)
(work with Yingwei Wang)

Blockchain is currently an important emerging technology. In general, blockchains were developed in order to secure digital transactions. It was first introduced as part of the cryptocurrency Bitcoin. Despite being used first for digital transaction security, it is not limited to this area. As a decentralized distributed ledger technology, blockchain has many applications such as securely storing patient's information, maintaining supply chain integrity, asset management, claim processing, etc. However, as a result of blockchain's design, each node in the blockchain system must continuously send and receive data as well as maintain a complete record of the blockchain data. As a result, the data size requirement for a node may exceed the capabilities of some devices, limiting its applicability to applications for mobile phones, tablets and personal laptops. As part of this research, a mobile application will be discussed, which implements a special kind of blockchain architecture allowing it to run on any android mobile phone or tablet, without the same storage limitations. This architecture, named Dewblock, is based on the principles of cloud-dew architecture, which come from the dew computing paradigm emerged in the post-cloud computing

era. The differences between Dewblocks and traditional blockchains will be discussed, including the two-tier communication protocol and the modified client model.

Designing of Hybrid Energy Efficient Systems using Reinforcement learning

Deepak Ramegowda (SFX)
(work with Man Lin)

Energy consumption is the most significant challenge in real-time embedded systems. With the increased use of portable devices like mobiles, laptops, and smartwatches, battery life has become the central focus in research and innovation. The modern processor employs a DVFS (Dynamic Voltage Frequency Scaling) for energy reduction. However, these techniques lack the intelligence to decide the optimal way of operating on platforms like Beagle Bone and Raspberry Pi that have captured DIY enthusiasts' attention and buddying hobbyists to save energy.

Hence introducing a novel form of reinforcement learning to unite intelligence to the system and perform computing would lead to energy-efficient design.

Approximating Conditionally Invariant Measures in Open Dynamical Systems

Adam Smith (UPEI)
(work with Shafiqul Islam)

Classically, the study of dynamical systems is concerned with points that stay in the system under transformation. Specifically, it is of great interest to find so called invariant measures that are particularly useful in describing the probabilistic behaviour of the system. Such systems are often used in modelling the real world, with a notable example being the (discrete) logistic population model. An open dynamical system is a system where points may leave the system and thus terminate their trajectory. Typically, these points escape through a 'hole' or multiple holes in the space. This notion of escaping from the system is also of interest for describing the real world, a popular example being a billiard table where balls (points) can escape by landing in one of six pockets (holes). In the pursuit of understanding a particular dynamical system, it is common to use numerical methods to approximate invariant measures. A numerical method employed in various areas of mathematics and physics is the maximum entropy method. We will take a look at using this method for finding (conditionally)

invariant measures for open dynamical systems and explore the mathematical significance of results.

VPRNN: Volume Preserving Recurrent Neural Networks

William Taylor-Melanson (UPEI)
(work with Andrew Godbout and Gordon MacDonald)

Recurrent neural networks (RNN) are a type of machine learning model used for sequence processing. Volume-preserving neural networks are neural networks composed solely of volume-preserving transformations in all layers, with the possible exception of the output layer. A volume-preserving transformation is one that does not change the Lebesgue volume of any measurable set in its domain. Linear transformations with this property are those with determinants having a magnitude of unity. Previous work by researchers at UPEI produced an explicit parameterization of these linear transformations using combinations of 2D rotations, specialized diagonal matrices, and permutation matrices. These transformations are used in the hope that volume preservation makes gradients less likely to explode or vanish during backpropagation, leading to easier training in very deep networks. Recurrent neural networks often struggle from the similar problem as very deep networks, leading to vanishing or exploding gradients for problems with long sequences. We present a new type of recurrent neural network that utilizes volume-preserving linear transformations in the recurrent step. These models admit unitary variants when the linear transformation has a restricted parameterization.

We show better performance than LSTM as well as other state-of-the-art unitary RNNs on several RNN benchmark tasks by utilizing both unitary and non-unitary variants of our VPRNNs in different problem settings.

GAT3D- Graph Attention Network for 3D object detection

Sumesh Thakur (SMU)
(work with Jiju Poovvancheri)

Convolutional neural networks used to detect objects in images rely on convolution operations. Although the convolution operation is effective, it requires a regular grid input. Unlike images, point clouds are usually sparse and unevenly spaced on a regular grid. Placing the point cloud on a regular grid will cause an uneven number of points in each cell. Traditionally, most 3D object detection algorithms focus on processing 3D point clouds using voxel grids or bird's eye

view (BEV). we propose a graph attention network to detect objects from a LiDAR point cloud. Towards this end, we encode the point cloud efficiently in a k-nearest neighbor's graph. We design a graph attention network, named GAT 3D, to predict the category and shape of the object that each vertex in the graph belongs to. We propose an attention-based feature aggregation technique in GNN for detecting objects in LiDAR scan. We first employ a distance-aware down sampling scheme that not only enhances the algorithmic performance but also retains maximum geometric features of objects even if they lie far from the sensor. In each layer of the GNN, apart from the linear transformation which maps the per node input features to the corresponding higher-level features, a per node masked attention by specifying different weights to different nodes in its first ring neighborhood is also performed. The masked attention implicitly accounts for the underlying neighborhood graph structure of every node and eliminates the need of costly matrix operations thereby improving the detection accuracy without compromising the performance. The experiments on KITTI dataset show that our method yields comparable results for 3D object detection.

Structural Priors for Vision-Based Segmentation of Web Pages

Saralin Zassman (MTA)
(work with Michael Cormier)

Web pages are a vital part of daily life for most people, but users with assistive needs often struggle to use the Web. Assistive interfaces such as screen readers often struggle to interpret the structure of a web page, and accessibility frameworks are often ignored or used improperly in practice. Using the visual appearance of a page, captured through an image of the rendered page, can provide implementation-independent parsing. The first stage is segmentation of the web page into a tree of semantically and visually coherent regions. The structure of the segmentation tree is complex, but design conventions in web pages can provide guidance. We will describe recent work towards an empirical prior probability distribution over web page segmentation trees, and its applicability to an existing Bayesian segmentation system.

Deadline-Aware Deep-Recurrent-Q-Network Governor for Smart Energy Saving

Ti Zhou (SFX)
(work with Man Lin)

Modern Artificial Intelligence has shown its ability to derive optimal control policies for a complex environment. Traditionally, to develop a well-functioning

system module, one must fully understand the underlying system's hardware features. In this work, we explore if it is possible to design an intelligent agent that can self-develop control policy to satisfy user requirements. We have designed and implemented a Deep Recurrent Q Net (DRQN) model controlling CPU frequency in the Linux kernel, saving energy while meeting the task's deadline set by users.