NEW YORK UNIVERSITY

COMPUTER SCIENCE DEPARTMENT

COURANT INSTITUTE OF MATHEMATICAL SCIENCES

Cloud Computing

Spring 2024 – Pr. Jean-Claude FRANCHITTI

(CSCI-GA.3033-026 - Thu. 7:10 - 9:10 pm)

Course Description:

This course is a capstone course that introduces the latest cloud multi-access computing platform services including support for Internet of Things (IoT), Artificial Intelligence (AI), Web3 decentralized platforms, and other technologies. These services may be leveraged to create industry 4.0 and 5.0 innovative Intelligent Autonomous Networked (IAN) business solutions and deploy them on hybrid multi-cloud platforms so they can operate at the network edge to minimize latency and power consumption. Students will learn about services provided on the top "Big Clouds", namely Amazon AWS Cloud, Google Cloud, Microsoft Azure Cloud, IBM Cloud, Oracle Cloud, Salesforce, and others. Related services provided on the Cloud and covered in the course include computing and hosting services, storage services, networking services, container management services (e.g., K8s), and shallow/deep machine learning services. Once familiar with Cloud services, Students will learn how to combine them to support cognitive computing and enable the creation of IAN solutions that cater to the practical needs of industry 4.0/5.0 and society 5.0 users in novel and creative ways. As they create these IAN solutions, students will learn about immersive Augmented/Virtual Reality (AR/VR) User Experience and Interfaces (UX/UI), hybrid decentralized computing and DApps, multi-access computing and Internet of Things (IoT), Mobility/5G networking and Software Defined Network (SDN) infrastructures, advanced robotics, and various key aspects of digital security that are critical to ensure the safety of modern software solutions. Throughout the course, students will develop their own Intelligent Autonomous Networked (IAN) business solutions using Cloud-based shallow/deep machine learning and other services and capabilities as noted above. Sample platforms and related application components will be provided as applicable to speed up the development and deployment process.

Course Sections and Topics:

Part 1: Introduction to Cloud Computing and Cloud Infrastructure Services

Session 1 – Part 1: Cloud Principles

The focus of this topic is to introduce students to the Cloud landscape and its service and deployment models.

Students will learn the following:

- Difference between Cloud Computing and On-Premise Computing
- How to access Big Clouds (such as the AWS Cloud, Google Cloud, Microsoft Azure Cloud, or IBM Cloud) via portals, APIs, and SDKs
- Cloud Computing definition and characteristics (e.g., elasticity, multi-tenancy, on-demand access, ubiquitous access, usage metering, self-service capability, SLA-monitoring)
- Cloud Computing and Service Oriented Architecture (SOA)
- Enterprise Cloud drivers and Adoption Trends
- Typical Cloud Enterprise workloads
- Cloud Service Models/Types (i.e., Public, Private, Hybrid, and Community)
- Cloud deployment models (i.e., IaaS, PaaS, SaaS, and BPaaS)
- Cloud Return on Investment (ROI) models
- Cloud Reference Architectures
- Cloud Standards (e.g., OSDIAPIs)
- Technology Providers vs. Cloud providers vs. Cloud vendors
- Planning Cloud transformations

Lab Assignment:

- Learn how to access the Cloud via Big Cloud vendors' websites, and their APIs/SDKs; install Python, Anaconda, and Jupyter to run lab notebooks as applicable

Session 1 – Part 2: Managing Data in the Cloud (Cloud Storage Services)

The focus of this topic is to introduce students to the various types of data storage systems that are available for use in the Cloud and illustrate how to interact with these services using a Cloud portal or directly with code.

Students will learn the following:

- Storage models and storage as a service
- Using Amazon Cloud Storage Services via Portal and APIs
- Using Microsoft Azure Cloud Storage Services via Portal and APIs
- Using Google Cloud Storage Services via Portal and APIs
- Using IBM Cloud Storage Services via Portal and APIs

- Using OpenStack Cloud Storage Services via Portal and APIS

Lab Assignment:

- Migrate a legacy storage structure to the various Big Clouds using the appropriate storage models and structures available on each individual platform.

Session 2 – Part 1: Cloud-Based Computing and Hosting Services

This part of the research program focuses on computing in the Cloud and understanding the spectrum of Cloud computing capabilities used to deploy single virtual machines or containers for simple interactive computing. Students will learn about computing as a service, service platforms, and using/managing virtual machines and docker containers

Students will learn the following:

- Virtualization in Cloud Computing Systems and Use of Hypervisors to Create Virtual Machines
- Computing as a Service
- Serverless Computing
- Using and Managing Virtual Machines on the Big Clouds
- Using and Managing Containers (Container Orchestration)
- Understanding and use of Docker containers and the Docker Hub
- Understanding and use of Google Kubernetes Engine and Alternative Container Management Solutions

Lab Assignment:

- Create a virtual machine on the various Big Clouds using both the Portals and the applicable python libraries; create/deploy a Docker container for the application created during Part 1 on Topic 2

Session 2 – Part 2: Scaling Cloud Deployments

The focus of this topic is to introduce students to techniques used to scale Cloud deployments via HPC, MapReduce, Graph Dataflow Execution, Agents and Microservices, etc.

Students will learn the following:

- Cloud Architectures and Infrastructure Design

- Dynamic Deployment of Virtual Clusters
- Parallel Computing in the Cloud
- SPMD and HPC-Style Parallelism
- Many Task Parallelism
- MapReduce and Bulk Synchronous Parallelism
- Graph Dataflow Execution and Spark
- Agents and Microservices
- HTCondor

Lab Assignment:

- Implement a client program that feeds data into a queue, to be consumed by microservices; use Spark to perform a simple MapReduce computation; use of a special set of commands to embed SQL in a Jupyter notebook directly; deploy Jupyter in a Spark cluster on an Amazon Elastic Map Reduce cluster, using the exploration of Wikipedia data as an illustrative example

Session 3: Cloud Networking Services and Service Platform Design

The focus of this topic is to introduce students to scalable Cloud computing capabilities used to deploy virtual machines to clusters of machines for data analytics or traditional high-performance computing. Students will also learn about the various platform services available on the Big Clouds.

Students will learn the following:

- Virtual Private Cloud Networking
- High-Performance, Scalable Load Balancing
- Cloud API Gateways
- Global Content Delivery Networks
- Cloud-Managed High-Performance Network Address Translation
- Network Edge Connectivity
- Reliable, Resilient, Low-Latency DNS Serving on the Cloud
- Network Performance and Availability Optimization on the Cloud
- Big Cloud Service Platforms Convergence and Service Offerings (Amazon AWS, Google GCP, Microsoft Azure, IBM Cloud, Force.com Cloud, Clouds at SGI, NASA, and CERN)

Lab Assignment:

- Setup a simple application on a VPC on the various Big Clouds and optimize its accessibility globally to facilitate testing in various remote locations.

Part 2: Cloud Platforms and Big Data Analytics and Machine/Deep Learning Cloud Services

Session 4 – Part 1: Cloud Big Data Analytics Services

This topic introduces students to Big Data Science and related challenges and prepares them to use applicable Cloud data analytics and related programming frameworks to implement machine learning applications. Students will learn about data collection, mining, and analytics on the Cloud. Students will also get practical exposure to supervised and unsupervised machine learning algorithms and model fitting.

Students will learn the following:

- Data Science and Big Data Characteristics
- Data Collection, Mining, and Analytics on Clouds
- Scalable Parallel Computing Over Large Clusters
- Data Analytics Frameworks (e.g., Hadoop programming with HDFS/MapReduce and YARN; Spark Core and Resilient Distributed Data Sets; Amazon Elastic MapReduce; Amazon Athena Analytics, Google Cloud Datalab)
- Example: Big Data Analytics for Healthcare Applications
- Example: Big Data Analytics for Social Media Applications

Lab Assignment:

- Use Google's Datalab to explore contagious disease records from the U.S. Centers for Disease Control, specifically looking at Rubella cases over a period of time; use Google's Datalab to examine weather station data and to identify an anomaly in one station's reporting

Session 4 – Part 2: Cloud Streaming Big Data Analytics Services

The focus of this topic is to introduce students to the activity of analyzing data coming from unbounded streams in real-time or near-real-time. This challenging activity is referred to as data stream analytics. An example of such data is that coming from instruments that control complex systems such as the sensors onboard an autonomous vehicle or an energy power grid. In that case, the data analysis is critical to driving the system and the value of the results diminishes rapidly as they get older.

Students will learn the following:

- Scientific Streams Examples
- Basic Design Challenges of Streaming Systems

- Data Stream Analytics Frameworks (e.g., Spark SQL and Streaming Programming; Amazon Kinesis and Firehose; Kinesis/Spark and the Array of Things; Azure Data Streaming; Kafka, Storm and Heron Streams; Google Dataflow and Apache Beam; Apache Flink)
- Example: Streaming Big Data Analytics for Social-Media Applications

Lab Assignment:

- Use Amazon Kinesis together with Spark to detect anomalies in data from Array of Things instrument streams

Session 5: Cloud Machine Learning Services

The focus of this topic is to introduce students Machine Learning, a sub-field of Artificial Intelligence, and to Cloud applications of Machine Learning. Recent breakthroughs in this area result from the convergence of the availability of big data, algorithmic advances, and faster computers.

Students will learn the following:

- Taxonomy of Machine Learning Methods
- Supervised Regression and Classification Methods (Linear Regression for Prediction or Forecasting, Decision Trees for Machine Learning, Bayesian Classifier, Support Vector Machines)
- Unsupervised Machine Learning Algorithms (Association Analysis, Clustering Methods Without Labels, Dimensionality Reduction, and Other Algorithms)
- Clustering and Dimensionality Reduction Methods (Cluster Analysis and K-Means Clustering, Dimensionality Reduction and Reinforcement Learning, Principal Component Analysis, Semi-Supervised Learning Methods)
- Model Development and Selection for Machine Learning (Model Over- and Under-fitting)
- Cloud-Based Machine Learning Frameworks (e.g., Azure HDInsight and Data Lake; Spark MLlib for Machine Learning and GraphX for Graph Processing; Azure Machine Learning Workspace; Amazon Machine Learning Platform)
- Example: Healthcare Problems and Machine Learning Related Applications

Lab Assignment:

- Use Azure's HDInsight and Spark to analyze food inspection records; implement a client that can be used to push data to a web service that implements a simple document classifier with the Azure ML tool

Session 6: Cloud Deep Learning Services (i.e., Cognitive Computing Differentiable Programming)

This part of the research program is a continuation of the previous topic. It delves into neuromorphic hardware and cognitive computing and prepares students to use Cloud programming frameworks to implement deep learning applications using differentiable programming. Students will learn about intelligent machines and deep learning networks.

Students will learn the following:

- Artificial Intelligence and Smart Machine Development
- Artificial Neural Networks (ANNs) and Forward/Backward Propagation for Deep Learning
- Taxonomy of Deep Learning Networks (e.g., Convolutional Neural Networks, Recurrent Neural Networks)
- Text and Image Recognition Using ANN and CNN
- Deep Learning of Other Brain Functions (e.g., Restricted Boltzmann Machines, Stacked AutoEncoder and Deep Belief Networks)
- Cognitive Computing and Neuromorphic processors
- IBM SyNAPSE, Cambricom NPU, Google's TPU, and related AI programs
- Classes of Deep Learning Networks (e.g., Recurrent Neural Networks – RNNs)
- Cloud-Based Deep Learning Frameworks (e.g., Amazon MXNet Virtual Machine Image; Google TensorFlow, Keras, DeepMind and graph analytics; Microsoft Cognitive Toolkit)
- Google TensorFlow for Neural Network Computing (Concepts and Innerworkings)
- Google TensorFlow System Applications for Deep Learning (Architecture, Handwritten Digit Recognition, Cognitive Services)
- Google's DeepMind Reinforcement Deep Learning Algorithm (AlphaGo and Other Programs)
- Predictive Software Libraries for Cognitive Applications (Keras and DIGITS for Deep Learning, and Graph Libraries for Graph-Parallel Computations in the Cloud and Community Detection in Social Networks)
- Example: Deep Learning Systems and Social Media Analytics

Lab Assignment:

- Deploy CNTK on a local machine; show how to reconstitute a recurrent neural network created by training CNTK with text from business news items, and then load and run the resulting model; use MXNet to train the Resnet-152 image

recognition model and load/run the train network to identify images from the web; use TensorFlow to built a simple logistic regression analyzer that can be used to make simple predictions of graduate school administration

Session 7: Cloud-Based Distributed Data Management Platforms and Other Cloud Topics

The focus of this topic is to first introduce students to Cloud capabilities that allow distributed data movement, data sharing, and credential and identity management. This is particularly useful as data may be created, consumed, and stored in a variety of locations (e.g., specialized scientific laboratories, national facilities, and institutional computer centers). Students will also learn as part of this topic that it is possible to build a basic Cloud from scratch using various open source software packages. Finally, this topic will touch upon managing (digital) security in the Cloud, which requires an understanding of Cloud architecture.

Students will learn the following:

- Challenge and Opportunities of Distributed Data
- Transferring and Sharing Data and Managing Identity and Credentials Using the Globus Platform
- Using Open Source Framework to Build a Cloud (e.g., Eucalyptus, OpenStack, VMWare)
    - Virtual Clustering in Private Clouds Using Eucalyptus
    - Private and Public Clouds Building Using OpenStack
    - Building Hybrid Clouds Using VMWare Virtualization Support
- Use Globus to Build a SaaS capability
- Security in the Cloud (i.e., Role-Based Access Control, Data Security, VMs and Containers Security, Cloud Software Services Security, Managing Trust in Clouds and Data Centers)

Lab Assignment:

- Analyze research data portal design patterns (portal vs. graph service) and build a remotely accessible service; stand up a single-cluster Cloud using Eucalyptus; stand up an HPC Cloud for scientific workload processing using OpenStack; leverage the Globus Genomics Bioinformatics System to provide a sample research data management service

Part 3: Hybrid Cloud and Decentralized Computing Platforms and Related Services

Session 8 – Part 1: Ethereum Basics

The focus of this topic is to introduce students to the basics of decentralized computing using Ethereum as a platform.

Students will learn the following:

- Ethereum Overview (Ethereum vs. Bitcoin, Components of a Blockchain, Ethereum History, Ethereum Development Stages, General Purpose Blockchain, Decentralized Applications (DApps), Web3, Ethereum Development Culture)
- Ether Currency Units
- Ethereum Wallets Overview
- Control and Responsibility
- Using MetaMask
- Externally Owned Accounts (EOAs) and Contracts

Lab Assignment:

- Implement a simple test ether Faucet, compile the Faucet contract, create the contract on the blockchain, and interact with the contract

Session 8 – Part 2: Ethereum Clients, Wallets, and Transactions

The focus of this topic is to teach students how to build decentralized applications clients on the Ethereum network, understand how to work with wallets, and perform transactions on the Ethereum Blockchain.

Students will learn the following:

- Ethereum Networks
- Running an Ethereum Client
- Synchronization of Ethereum-Based Blockchains
- Remote Ethereum Clients
- Cryptography Basics (Private/Public Keys, Cryptographic Hash Functions, Ethereum Addresses)
- Wallet Technology Overview
- Wallet Best Practices
- Structure of Blockchain Transactions
- Transactions Nonce, Gas, Recipient, Value and Data
- Special Transaction: Contract Creation
- Digital Signatures
- Signature Prefix Value and Public Key Recovery

- Offline Signing
- Transaction Propagation
- Recording in the Blockchain
- Multiple-Signature (Multisig) Transactions

Lab Assignment:

- Deploy a simple Ethereum client application provided to perform transactions on the blockchain using MetaMask and the default wallet.

Session 8 – Part 3: Smart Contracts, Solidity, and Vyper

The focus of this topic is to teach students how to implement Smart Contracts using Solidity and Vyper on the Ethereum Blockchain.

Students will learn the following:

- Overview of Smart Contracts
- Ethereum High-Level Languages
- Building a Smart Contract with Solidity
- Ethereum Contract ABI
- Programming with Solidity
- Gas Considerations
- Vulnerabilities and Vyper
- Vyper vs. Solidity
- Decorators and Function/Variable Ordering
- Compilation and Management of Overflow Errors
- Reading and Writing Data Using Vyper

Lab Assignment:

- Augment the simple Ethereum client application provided to implement Smart Contracts using Solidity and Vyper.

Session 9 – Part 1: Decentralized Applications (DApps)

The focus of this topic is to introduce students to Decentralized Applications (DApps).

Students will learn the following:

- DApps Overview
- Using Swarm

- The Ethereum Name Service (ENS)
- Differences Between Apps and DApps

Lab Assignment:

- Implement the back-end Smart Contracts and front-end user interface of a basic DApp Auction example, further decentralize the sample Auction DApp, store the Auction DApp on Swarm, and use the ENS to access the DApp.

Session 9 – Part 2: Ethereum Alternatives and Other Topics

The focus of this topic is to introduce students to Ethereum alternative platforms as well as cover additional topics related to DApps.

Students will learn the following:

- BlockStack vs. Ethereum Blockchain
- HashGraph vs. Ethereum Blockchain
- Hybrid App/DApp Platforms (HDApps)
- Smart Contract Security Detailed
- Tokens, Related Standards, and Usage on Ethereum
- Tokens and Security Tokens, ICOs, and SCOs
- Oracles
- Consensus
- Ethereum Virtual Machine
- Upcoming Blockchain Trends (Blockchain As a Service, Federated Blockchain, Ricardian Contracts, Blockchains Interoperability, Stable Coins)

Lab Assignment:

- Review the hybrid App/DApp platform provided and augment the App/Dapp (HDApp) that is implemented on top of it to add both Smart Contract(s), and React User Interface(s) extensions

Part 4: Intelligent Autonomous Network (IAN) Solutions and Additional Key Services

Session 10: Cognitive Process Modeling and AI

In this part of the research program, students will learn about cognitive process modeling and the use of AI to support such modeling.

Students will learn the following:

- Human Intelligence and Cognition
- Computational Approaches Used to Understand Human Intelligence and Cognition
- Computational Cognitive Modeling to Understand Human Behavioral Data (Mind and Brain More Generally by Building Computational Models of the Cognitive Processes that Produce the Data leveraging ANNs/Deep Learning, Bayesian/Structured Probabilistic Models, Symbolic and Logical Systems, and Reinforcement Learning)
- Goals, Philosophy, and Technical Concepts Behind Computational Cognitive Modeling (Model Fitting and Evaluation)
- Psychological Applications of Computational Cognitive Modeling Approaches (e.g., Learning, Memory, Decision-Making, Language, Categorization Reasoning, and Problem Solving)
- How Cognitive Science Can Inform Research in Machine Learning and AI
- How to Fit and Evaluate Cognitive Models to Help Understand Behavioral Data
- How to Build IAN Business Solutions that Improve their Overall Performance Proactively to Meet Upcoming End-User Needs and Expectations

Lab Assignment:

- Leverage the hybrid App/DApp platform provided and augment the sample IAN App/DApp (HDApp) that is implemented on top of it to collect and analyze both user as well as App/DApp behavioral data that can be used within the IAN application to improve overall performance proactively

Session 11: Immersive Augmented/Virtual Reality (AR/VR) User Experience and Interfaces (UX/UI)

The focus of this topic is to introduce students to AR/VR technology used to build immersive user experience for IAN business solutions.

Students will learn the following:

- Augmented, Mediated/Mixed, and Virtual Realities (AR, MR, VR)
- AR/VR Innerworkings and Available Platforms
- Holograms and Holographic Display Overview
- Microsoft Hololens 2 AR Development Environment Setup (Visual Studio 2017, Windows 10 SDK, Hololens Emulator and Holographic Templates, Unity Game Engine, and Vuforia as optional)
- Microsoft Hololens 2 API Programming using C#, C++, or JavaScript
- Buildwagon JavaScript Development Platform for Microsoft Hololens 2

- Holoportation
- Microsoft Hololens 2 Tutorials, Academy, Wiki, Forum, and Hands On

Lab Assignment:

- Stand up a Microsoft Hololens 2 Development Environment to extend the App/Dapp platform and sample IAN App/DApp (HDApp) application provided with immersive AR UX capabilities to improve the overall user experience of the application

Session 12 – Part 1: Edge and fog computing and Internet of Things (IoT)

The focus of this topic is to introduce students to Edge and Fog and Internet of Things Technology. Fog computing is an approach that enables computing services to reside at the edge of network and closer to sensors and/or end-users. It achieves better Quality of Service (via latency reduction and backbone bandwidth savings) and edge analytics/stream mining to optimize user-experience.

Students will learn the following:

- Sensing Technologies for Internet of Things
- Local and Global Positioning Systems (GPS)
- IoT Interactions with GPS, Clouds, and Smart Machines
- Radio Frequency Identification (RFID)
- Sensors, Wireless Sensor Networks and GPS Systems
- Cognitive Computing Technologies and Prototype Systems
- Cloud-Based Radio Access Network (RAN) for Building Mobile Networks
- IoT Interaction Frameworks with Clouds and Devices
- Example: IoT-Based Healthcare Systems and Applications

Lab Assignment:

- Extend the App/Dapp platform and sample IAN App/DApp (HDApp) application provided to improve Quality of Service and edge analytics/stream mining and optimize the overall user-experience of the application

Session 12 - Part 2: Mobility/5G Networking & Software Defined Technical Infrastructures

The focus of this topic is to introduce students to mobility/5D networking technology and Software Defined Secure Technical Infrastructure. These technologies are meant to drastically improve QoS and overall user experience when using IAN mobile applications .

Students will learn the following:

- Wireless Internet and Mobile Cloud Computing (Mobile Devices, Internet Edge Networks; Wi-Fi, Bluetooth, and Wireless Sensor Networks; Cloudlet Mesh for Mobile Cloud Computing; Mobile Clouds and Colocation Clouds)
- Software-Defined Networks
- Software-Defined Storage
- Software-Defined Compute
- Automating Cloud Applications' Security Scanning
- Optimizing the Protection of Cloud Services Against DoS and Web Attacks
- Leveraging Network Telemetry to Optimize the Security of Cloud Services

Lab Assignment:

- Extend the App/Dapp platform and sample IAN App/DApp (HDApp) application provided based on the technology covered in this topic to further improve Quality of Service and edge analytics/stream mining and optimize the overall user-experience of the application

Session 13: Optimizing High-Demand Business Solutions Using IAN Applications' Transmutations

The focus of this topic is to introduce students to real-time IAN applications transmutations to enable proactive improvements via discovery and dynamic addition of functional features in order to meet upcoming end-user needs and expectations.

Students will learn the following:

- Intelligent Autonomous Optimization of Networked Business Solutions in a Cloud and Decentralized Architecture Context
- (Cloud) Mashup Architecture for Extensibility Support and Quality of Service (QoS) (e.g., Agility, Scalability, Performance, and Related Metrics)
- Mobile Clouds and MultiCloud/Inter-Cloud Mashup Service Architecture
- Skyline Discovery of (Mashup) Services
- Dynamic Composition of (Mashup) Services
- IAN and Social Media Applications Analysis (Leveraging Social Media Applications APIs and Social Graph Properties, Performing Social Graph Analysis on Smart Clouds via Deep Reinforcement Learning)
- Example: IAN Emotion-Control Healthcare Applications

Lab Assignment:

- Extend the App/Dapp platform and sample IAN App/DApp (HDApp) application with the ability to discover and dynamic add functional features in order to meet upcoming end-user needs and expectations and further optimize the overall IAN application user-experience

Additional Topics:

As time allows, the research program will delve into further IAN applications optimization capabilities beyond applicability and basic suitability such as performance, security, data privacy, reliability, and other 'ilities'.

Projects:

Throughout the course, students will develop their own Intelligent Autonomous Networked (IAN) business solution using Cloud-based machine and deep learning and other services and capabilities as noted above. Sample platform and related application components will be provided as applicable to speed up the development process.

Reading Materials:

The recommended textbooks and references for the seminar and research program are as follows:

Parts 1 and 2:

- Cloud Computing for Science and Engineering, by Ian Foster and Dennis B. Gannon (ISBN: 978-0-262-03724-2)
- Cloud Computing for Machine Learning and Cognitive Applications, by Kai Hwang (ISBN: 978-0-262-03641-2)
- Big Clouds websites (e.g., azure.microsoft.com, aws.amazon.com, cloud.google.com, cloud.ibm.com, force.com)
- Instructor-provided slide sets, handouts, and research papers

Part 3:

- Mastering Ethereum, Building Smart Contracts and DApps by Andreas M. Antonopoulos and Dr. Gavin Wood (ISBN: 978-1-491-97194-9)
- Various websites (e.g., ethereum.org, blockchainhub.net, blockstack.org)
- Instructor-provided slide sets, handouts, and research papers

Part 4:

- Instructor-provided IAN book manuscript, slide sets, handouts, and research papers
- Online TED talks

- Various websites (research.ibm.com, Big Clouds websites)

Prerequisites:

Students should have basic computer knowledge and familiarity with one or more scripting or programming language (e.g., Python, JavaScript, Java). They should come to class with a laptop and will need to research the various topics and submit a two-page paper (e.g., lab report) before the beginning of each session.

Hands-On Approach:

Students will attend lectures and will be given access to material and data sets to perform hands-on exercises during classes. They will also need to document their exercises solutions, complete practical homework assignments, and work on an ongoing team project between classes. Students will need to successfully complete a midterm exam and a final exam (take home) project to pass the course. The final exam will be a combination of an ongoing course project submission and a final presentation of the project. Teams will be allocated to work on the ongoing course project early on during the course. With the help of the instructor, each team will identify an IAN business solution to implement as part of the project. Teams will use an agile solution development lifecycle methodology to work on their projects and project development will start at the beginning of the second week of class.