UNIVERSITY OF MIAMI

"Deep Analytics for Massive Data: Roadmap Towards a Transformative Knowledge Grammar Using Artificial Intelligence, Machine Learning, and Data Science" (SYLLABUS for 3-Week Summer Faculty Workshop | June 7-25, 2021)

Recent disruptive progresses in artificial intelligence, machine learning, and data science are redefining the methodology and narrative of scientific investigation and knowledge presentation. Currently academic research adaptation to this new "knowledge grammar" is just starting and this transformative roadmap leads to many promising frontiers. In this workshop, we will present a timely update of the innovations, success stories, and emerging technical trends. Furthermore, we will "design success": transformative methodologies, protocols, and reference projects directly translate into cutting-edge knowledge while forging synergistic collaborations.

Different from on-the-market workshops of artificial intelligence, machine learning, and data science (AI/ML/DS), our presentation focuses more toward academic adaptation instead of industrial adaptation. Most on-the-market workshops are tailored towards industrial developer for harvesting new technology for product prototyping in a few "trendy" areas (web services, intelligent/connected devices, etc.). This industrial-focused approach usually limits the presentation to a curated garden of technologies, a "hype dictionary" without emphasizing the connections between these technologies, their "ecological" context, and the rationale for "shopping" among choices. Our presentation will provide a roadmap for academic adaptation that encompasses a wider aperture compared to industrial adaptation. Beyond harvesting existing technologies, we will focus more toward the creation of new knowledge and place our talents at strategic frontiers, with in-depth presentation of research literature and reference research projects that defines a transformative knowledge horizon for fundamental innovations.

This introductory course aims at a comprehensive but accessible coverage of core concepts and technology for AI/ML/DS, without prerequisite while steering perspective talents to potential technique areas. Each lecture is 60 minutes. Each lab is 60 minutes introduction session plus 2-hour workload for reading and implementation. We use an application-driven organization that covers the technologies in several application areas. Then we cover mathematical skills and programming tools when needed. This approach allows us to cover exciting technologies earlier in the course, while providing rigorous foundations and skills needed for academic research and publications.

Part I: Concepts

Introduce concepts and backgrounds for AI/ML/DS.

Lecture 1: AI/ML/DS Roadmap

History, current progress, fundamental concepts of intelligence and computing.

Lecture 2: AI/ML/DS Academic Research

Academic research on AI technology, applications to many academic areas, requirements for a data scientist in academic environment.

Lecture 3: AI/ML/DS Industry Adoption Approaches

Industrial research and development, products, economic models, skills for industrial data scientists.

Part II: Understanding Natural Language

Natural language applications on AI. We cover both pre-deep-learning techniques from natural language processing and computational linguistics and the frameworks based on deep learning systems.

Lecture 4: Natural Language Concepts: Representations and Modeling

Modeling natural language, grammar, semantics, speech signal processing, and machine translation.

Lecture 5: Mathematical Foundations for Natural Language Processing
Matrix computation and optimization, models for text embedding

Lecture 6: Deep Neural Networks
Neural networks, deep neural net architectures, learning and inference processes.

Lecture 7: Recurrent Neural Network Architectures
Recurrent architecture, LSTM/GRU, classification tasks.

Lecture 8: Encoder-Decoder Networks and Machine Translation
Sequence-to-sequence models, generative tasks, dialogue systems, machine translation systems.

Lecture 9: Attention Mechanism
Visual attention, temporal attention, transformer architecture.

LAB 1: Gain Insights from Customer Reviews
Understanding customer sentiment from online reviews using text classification frameworks.

LAB 2: Creating an AI Dialogue System
Construct a dialogue system based on transformer architecture.

Part III: Computer Vision

Lecture 10: Image Representation and Processing
Data structure, signal processing, image feature extraction and analysis.

Lecture 11: Mathematical Foundations for Computer Vision
Vector calculus, matrix computation, similarity qualification.

Lecture 12: Convolutional Neural Network Architectures
Layers of Convolutional Neural Nets, common architectures, comparisons of models, transfer learning, classification and localization tasks.

Lecture 13: Understanding Multimedia: Multimodal Modeling
Video analysis, audio and music features, integration of multiple modalities, collecting audience responses.

LAB 3: Analyze, Classify, & Detect Objects
Annotate the objects in an image, classify similar objects (challenging classification tasks such as similar plant types), locate object within a picture.

Lecture 14: Deep Neural Networks: Information Dynamics
Modeling and analysis techniques for deep neural net, training and inference procedures.

Lecture 15: Deep Neural Networks: Advanced Architectures and Neural Architecture Search
Advanced architectural features, optimization-based neural architecture selection, performance evaluation metrics.

LAB 4: Classifying Images and Automatic Feature Engineering

Classify user-collected photos, automatically highlight image features (image areas important for classification).

Part IV: Robotics, Games, and Creativity

Lecture 16: Perception, Planning, Decision, and Control
Introducing key concepts and applications.

Lecture 17: Mathematical Foundations for Robotic Control
Uncertainty, optimal control, and decision.

Lecture 18: How to Play Go and Electronic Games: Deep Reinforcement Learning
Policy search, neural net policies, policy gradients, Markov decision processes, and deep Q-learning

Lecture 19: AI as Artists Generative Adversarial Networks
Training simple GANs, data generators, discriminators, Nash equilibrium, model optimization, deep convolutional GANs, style transfer.

LAB 5: Virtual Space Flight
Plan and simulate a robotic space exploration task.

LAB 6: Play Go with Champions
A simple algorithm for a Go or chess playing virtual agent.

Part V: Tabular Data Analytics and Decision Support
This part covers AI impact on tabular data (spreadsheet data).

Lecture 20: Exploration, Preparation, Representation, and Transformation
Data organization, annotation, handling missing values, preprocessing techniques.

Lecture 21: Data Visualization and Presentation
Visualization tools, storytelling with data.

Lecture 22: Exploratory Data Analysis and Probability Density Estimation
Tools for exploring univariate and multi-dimensional data, principle component analysis and other dimensionality reduction tools, kernel density estimation, finite mixtures.

Lecture 23: Supervised Learning
Naive Bayes classifiers, Bayesian decision theory, classification tree, random forest and boosting, nearest neighbor, support vector machines, combining classifiers, linear models, sparse modeling.

Lecture 24: Performance Metrics, Interpreting Big Data, and More
Accuracy, precision, recall, receiver operating characteristic curve, statistical significance/confidence, Bonferroni's principle, false discovery rates, large-scale hypothesis testing.

Lecture 25: Unsupervised Learning and More
Similarity distances, hierarchical clustering, k-means, finite mixture, evaluating clustering results.

Lecture 26: Massive Data Methodologies
MapReduce and Spark, similarity-preserving summaries of sets, locality-sensitive hashing, large-scale machine learning, mining data streams.

LAB 7: Exploratory Data Analysis
Dimensionality reduction for high-dimensional data, kernel density estimation for time series, curating analysis results.

LAB 8: Predicting Customer Usage Pattern
Tracking customer behavior based on high-dimension features.

Part VI: Programming Supercomputers

Lecture 27: Scaling Up Towards Supercomputers
Heterogeneous parallel computing, architecture of a modern CPU/GPU, architecture of supercomputers, parallel programming languages and models, data parallel computing, scalable parallel execution.

Lecture 28: Memory and Data Locality
Memory access efficiency, CUDA memory types, tiling for reduced memory traffic, boundary checks, memory as a limiting factor to parallelism.

Lecture 29: Performance Optimization
Memory bandwidth, memory parallelism, warps and SIMD Hardware, dynamic partitioning of resources, thread granularity.

Lecture 30: Parallel Patterns
Stencil computation, speed and work efficiency, atomic operations and privatization, data compression and regularization, tiling with dynamic input data identification, graph search.

Lecture 31: CUDA Dynamic Parallelism
Dynamic parallelism, memory data visibility, configurations and memory management, synchronization, streams, and events.

Lecture 32: Programming a Heterogeneous Computing Cluster
Message passing interface, overlapping computation and communication, message passing interface collective communication, CUDA-aware message passing interface

Lecture 33: Parallel Programming with OpenACC
OpenACC execution model, directive format, comparing OpenACC and CUDA, Interoperability with CUDA and libraries.

Lecture 34: More on CUDA and Graphics Processing Unit
Model of host/device interaction, kernel execution control, memory bandwidth and compute throughput, programming environment.

LAB 9: Big Data Prediction
Machine learning using parallel architectures, performing real-time prediction based on high-dimensional "massive" dataset.

LAB 10: Simulating Dynamic Systems
Construct a large-scale dynamic system using the interaction models of the components. Detect large-scale 'emerging" properties from collective behaviors.

Lecture 35: Ethics of AI/ML/DS
Policies, intellectual property laws, philosophies.