Jitendra Upadhyay

Contact Details: +91 9606189685, +65 90524938
Skype: Jitendraupadhyay10@gmail.com

Profile

- Master of Business Analytics post graduate with directly relevant financial services ML background
- Delivered solid engineering solutions across many areas of the Data and AI spectrum. Experienced from each discipline to be able to accomplish end-to-end Data Science solutions. From developing Machine Learning algorithms to creating complex, Big Data-ready ETL pipelines to deploying Machine Learning as a Service. Jitendra has taken a project from idea inception to API and deployed to cloud.
- Deep learning: **BERT**, Reformers, ULMFIT, FLAIR, Keras Model, TensorFlow 2.0, Pytorch.
- Artificial intelligence: artificial neural networks, rules-based models, fuzzy logic controllers, genetic algorithm, and ant-colony optimization. RNN, CNN
- ML Probabilistic programming in Figaro / Scala & deep probabilistic programming in Edward / TensorFlow, ABBYY, Tesseract4, OpenCV
- Reinforcement learning: TD Lambda, Policy Gradients, DQN, and A3C in TensorFlow
- Distributed programming in MapReduce / **Java** and Python, TensorFlow, H2O (mrjob), **Spark / Scala** and **Python**, TensorFlow, H2O, Akka / Scala and H2O; Designed CDH Hadoop cluster with up to 30 nodes; Architected and Designed Hadoop cluster consisting of Edge nodes, Master node and Data Nodes
- Open source libraries & toolkits: mrjob, Apache Mahout, MLlib, Apache Solr / Lucene, WEKA's API, NumPy, SciPy, PyMC, Scala Breeze, Figaro / Scala, Edward, NLTK, spaCy, MALLET, Mulan, Gensim + word2vec, scikit-learn + nolearn, milk, SVM-Light / PySVMLight, Pycluster *, Yolo4, Imagenet, Alexnet, ELECTRA*
- Platforms: scientific computing (**Spark, Keras & TensorFlow**, WEKA / MEKA, IPython, TensorFlow, H2O / Jupyter, OpenAI Gym, NetLogo), cloud computing (Google App Engine, Amazon EC2), Windows, Linux (Fedora, Ubuntu), MongoDb, **Cloudera Hadoop,** *Lambda, Kinesis, Google Biq Query , Data Flow, Cloud CDN.*
- Data Science/Data Engineering tooling: AWS Sage maker, **Apache Airflow**, ML Flow, AWS Batch. Devops / DataOps,
- CI: **Docker,** GitHub, Kubernetes, GKE.
- NLP / Deep NLP & text mining: tagging (based on a trigram HMM), syntactic parsing (based on a PCFG), feature engineering and dimensionality reduction, multi-label classification, word sense disambiguation, Twitter hashtag decomposition, relevance engine, topic modeling; sentiment analysis, contextual text mining.
- Data mining & machine learning: Recommendation & personalization (based on real-time reinforcement learning) systems model free and model-based environments.
- Using the state-of-the-art in computer vision and deep learning, including 3D and video applications.

- Develop and implement state-of-the-art computer vision algorithms for object detection, classification, segmentation, etc., including advanced medical diagnostic use cases.

- Translate business requirements into quick prototypes and work with clients directly to uncover operational objectives and advance data collection.

Work Experience

Cartrack Singapore

AI Engineer / Senior Data Scientist Jan 2020 – Present

- Developing models to create intelligence and efficiency solution provision to over a million subscribers, many of whom are large fleets. Cartrack is creating revolutionary platforms that will change the way that people interact with their vehicles.
- Project: Anti-fraud System for Car Insurance Claim Based on Visual Evidence
  - Developing the underlying model methodology and the car damage detector web application
  - Uses RCNN architecture to recognize car damage from a given image and depict the underlying targeted classes to generate a report for auto insurance claim purpose. This development attempts to classify car damages into a few specific classes.
    - Multi-sensor Fusion in Automated Driving: sensors such as cameras, LIDAR, radar, sonar, GPS, IMU.
    - o Used SLAM, 3D reconstruction, and calibration.
    - o Experience with image-space algorithms such as edge detection, segmentation, optical flow, scene flow, and image decomposition.
    - o Knowledge in deep learning 3D perception
- Project: Route optimization using reinforcement learning from Google Image data
- Technologies: Pytorch, TensorFlow, GCP, Amazon Sage maker, Power BI, Kubernetes, OpenCV, PCL, ROS), CUDA/OpenCL,

IBM GSL Analytics June 2019 – Dec 2020

My team has developed lots of advanced techniques and products for IBM Watson AI. Our mission includes designing new deep learning techniques and novel efficient product features to improve IBM Watson services and accelerate IBM technique transformation. The goal is to facilitate our internal and external business partners. Another mission is to push the boundary of state-of-the-art AI techniques for IBM and the public deep learning community. Our research and products cover natural language understanding, computer vision, sensor signal understanding, medical and digital health, etc.

Computer vision and Deep learning related projects.
-Fine-grained object classification
-Deep model compression and acceleration
-Mobile scalable deep learning model design.
*Mobile based deep model optimization with efficient power and latency.
*Improved the running speed of deep module on mobile by 9x.
*Deliver the scalable deep learning model with 6Mb (fine-grained object classification) with more than 95% accuracy (field test in the wild conducted by the QA group)
-Submitted three patents and eight papers in the past one year. These topics covered deep model compression and acceleration, object classification, multi-task learning, deep learning based mobile privacy and latency optimization, face image analysis, data mining, etc.

Tabsquare.AI, Singapore FreeLancer

ML Senior Engineer Mar 2019 – June 2019

- Built a Goal-Oriented Chatbot with Deep Reinforcement Learning. Developed the system comprising dialogue system for a GO chatbot using reinforcement learning is split into 3 main parts: The Dialogue Manager (DM), Natural Language Understanding (NLU) unit and Natural Language Generator (NLG) unit. The DM is a Dialogue State Tracker (DST) with a policy for the agent itself, which is represented by a neural network distant-supervision based technique for highly accurate (89%) Relation Extraction for factoid Question-Answering using Sequential Neural Networks.
- Analyzed and model structured data using advanced statistical methods. Created advanced demand forecasting functions powered by machine learning. This is supported by analytics and collaboration tools allowing operators to see the present and plan for purchasing and staff deployment.
- Developed a predictive, multi-layered algorithm to calculate a daily baseline forecast, which can be broken down into periods as short as 15 minutes by sales item. The forecast considers everything that affects demand: data from the same day last year, last month and last week, notable days, recent trends, last year's trends, public holidays, local and national events and the weather. Managers can refine the forecast further with local information, like roadwork outside the restaurant.
- Programming languages: Python, TensorFlow, H2O, Elastic Search, Julia, Hadoop, Java, Apache Spark, Scala, SQL and Shell scripting. Hands on with OCaml, Cyton/C, TensorFlow, Hadoop HDFS, AWS S3 and R. PyTorch, Keras, OpenCV, Caffe/Caffe2, H2O and Tensor Flow API.
- Experience with mathematical and statistical Python, TensorFlow, H2O, Elastic Search, Julia, Hadoop libraries such as pandas, scikit-learn, NumPy and SciPy, TensorFlow, Keras, Owl, and software: MATLAB and Mathematica.

Bank of America

Senior Data Scientist Jan 2018 – June 2019

- Responsible to support a DOE platform to run controlled online experiments at scale. Implement hypothesis testing techniques to estimate sample size, A/B testing, Multi-variate testing, Heterogeneous treatment effect analysis.
- Worked in the training the goal oriented chatbot ERICA for Merrill Lynch users created DQN and Dialog state tracker for the same using Reinforcement learning and Deep Learning.
- Build a sophisticated auto-ML pipeline to automate ML models. This tool is leveraged by data scientists and non-technical users. It saves 30-40% of data scientist development time-efforts.
- Worked extensively on applying NLP techniques to capture client's complaints from the chat transcripts. The model is used by the complaint supervision team. It helped reducing more than 50% of the manual effort time.
- Built and implemented predictive models by utilizing various databases to forecast volume change for complaints across the enterprise to better allocate customer service resource to meet client demand
- Pioneered the use of Machine Learning and Data Mining methods to automatically categorize complaints from agent-customer conversations, customer profiles and other variables

- Deduced topics and sentiments from unstructured social media contents using Machine Learning and Text Mining techniques to predict future types of complaints so preventive campaigns can be installed
- Detected trends and patterns hidden in unstructured customer complaints using SAS Text Mining models to automatically provide interactive visualization into topics, sentiments and clusters related to satisfactory issues
- Supported the introduction of SAS Contextual Analysis to BAC and developed/validated the rules to categorize customer complaints according to Regulators, Executives, Legal and Compliance's requirements for easy focus on complaint details
- Incorporated advanced data visualizations into Analytics & Reporting to allow senior management to review vast amounts of information in more efficient and constructive ways
- Developed algorithms for Augmented Reality, Image Registration, Object Tracking, Object Removal, HDR Imaging
- Researched / development of computer vision algorithms centered on object recognition (Python, TensorFlow, H2O, Elastic Search, Julia, Hadoop) Computer Vision Architect
- Improved customer base by profiling companies based on the 300M banking transactions, which involved clustering similar businesses transaction under one entity by using Machine Learning algorithms and further improved the model by reinforcement learning.

  • It enhanced selection of target audience, reduced manual intervention, and decreased false positive cases.

  • II) Designed and implemented machine leaning classification methods to classify 300M business transactions based on category and channel type into 60K business entities.

  * Achieved 95% accuracy in classifying transaction and clearly segregated wrongly classified transactions which reduced the manual intervention upto 80%.

  • III) Feature engineered audit evaluation metric records for record classification.

  IV) Worked on Named Entity Recognition.

  * Hands on experience on Python, TensorFlow, H2O, Elastic Search, Julia, Hadoop ML libraries (Keras, Scikit, H2O) with Kmeans, Phonetic algorithms, Apriori, String Distance algorithms, Random Forest, Naive Bayes, SVM, Regression, Deep Neural Network

Seagate Technologies India

Senior Data Scientist Jul 2015 – Dec 2017

- Led machine learning initiatives in anomaly detection, predictive maintenance & predictive intelligence for improving yields with very little decision latency during production.
- Architect for new machine learning model scoring service with ability to serve models built in several different languages and frameworks.
- Accomplishing Smart Factory 4.0 with deep learning framework for the real-time automated inspection of wafer images; presented at the 2017 Advanced Process Control (APC) conference.
- Designed analytics dashboards & reports on an ad-hoc basis for senior management to visualize and take actions key critical product metrics (KPI), process drifts & machine health.

- Estimation of Process Control Monitoring (PCM) test parameters using Machine Learning (SVM) with a reduction in average error of 53% to 18% across the wafer.
- Identified defect patterns in Hard Drive components using Machine learning techniques like Convolutional Neural Networks to improve the accuracy, reduce the errors and identifying the root causes of those faulty patterns during the manufacturing process.
- Implemented a Python, TensorFlow, H2O, Elastic Search, Julia, Hadoop variation of a learning algorithm for Generalized Additive Models. Contributed to key algorithms to generate model insights and auditability.
- Performed statistical analyses to quantitatively assess the upper airway in infants and children with subglottic stenosis
- Conducted advanced statistical analyses to determine if routine computed tomography imaging is necessary in the workup for children with connexin-related sensorineural hearing loss
- Studied the structure of higher-order tensor by extending the conventional SVD/PCA
- Tested the location effect and seasonal effect by t-test and ANOVA model within R
- Supported Data Scientists with ad-hoc and production algorithms for feature analysis and selection. Provided dashboards and automated reports for business stakeholders.

Tata Consulting Services, London, Palo Alto CA & India

Deep Learning Engineer / Infrastructure Quality Engineer to Citibank April 2010 – Jul 2015

- Implement rapid-time-to-market and applied innovation for lines of businesses and IT involving Life and Annuities, Retirement Planning Services, Group Protection products and services
- Leading and performing both strategic and hands on work on predicting strong drivers for successful sales and effective marketing strategies through machine and deep learning
- Implemented advanced Lifetime Value (Pareto/Beta-Geometric/Gamma-Gamma) and machine learning models to predict top advisors in our distribution network who can close sales with the highest probability matched against the opportunities in the market they serve
- Analyzed TB-sized, disparate customer-dataset and implemented new propensity model pipeline using Apache Spark, surfacing previously unknown churn indicators.
- Solidied and scaled end-to-end PySpark ETL-machine learning pipeline, resulting in a 5x increase in handled data-scale and 5x decrease of training time.
- Reduced productization times of new machine learning product features by 3x by creating new featurization prototypes in quick iterations with product and data-science teams.
- Deployed, debugged and maintained complex, distributed software stacks, containing Apache Spark, Hadoop HDFS and I Python, TensorFlow, H2O, Elastic Search, Julia, Hadoop Notebook servers, on cloud-based AWS systems. Optimized the stacks for best computational performance and stability.
- Architect of new machine learning model scoring service with ability to serve models developed in several different languages and frameworks.

- Implemented Python, TensorFlow, H2O, Elastic Search, Julia, Hadoop variants of various learning algorithms, such as Generalized Additive Models and Constrained Linear Models to generate model insights and auditability for regulatory compliance.
- Supported Data Scientists with ad-hoc and production algorithms for feature analysis and selection. Provided dashboards and automated reports for business stakeholders.
- Developed and deployed several models for credit underwriting, including models for new products.
- Analyzed and integrated new data sources into production systems to increase data redundancy
- As a Senior Data and Analytics expert dealing with several BI/ETL technologies such as QlikView, Qlik Sense, Power BI, MicroStrategy, SAP EPM, Talend.
- Support various lines of business with innovative transaction analysis:
  - Credit Card Collections: Optimized collections calling based on customer deposit data to provide lift of millions of dollars in collections productivity.
  - Mass Affluent: Identify and size transactions indicating a likely mass affluent customer. Identify customers transacting with brokerage competitors.
  - Commercial Credit Risk: Identify and quantify retail sales by merchant to provide timely information on revenue performance
- Structured quantitative investment strategies by researching dynamic asset allocation using optimization and statistics/econometrics techniques
  - Designed hedging strategies for pension clients, insurance companies, and financial institutions via analysis of correlations among a wide range of equities' historical price time series data
  - Examined the relationship between the USD 10-year vanilla interest rate swap and the Brazilian Ibovespa Index and designed an investment strategy by buying TLT Put options and finance them by selling IBOV Call options to position investors for a higher growth/higher yield environment in the U.S.
  - Priced derivatives such as barrier options, credit default swaps and risk reversals using Python (Quartz)
  - Investigated the contribution of price change due to earnings release to the change of stock volatility surface via analysis of covariance (ANCOVA) model
  - Adopted/utilized quantitative framework to assist live derivatives trading in daily analytical desk support
- Oversaw the build and release processes ensuring the integrity of changes released into the QA and Production environments. Managed the software development lifecycle and work with users and business analysts to implement functional specifications by delivering key solutions
- Faced off to EMEA and Asia Pacific Product Controllers by acting as the primary point of contact within the Repo IT team for enhancement requests and issue investigation
- Evaluated new technologies/components mandated for strategic use within the firm by building out Proof of Concepts and oversee their adoption following management approval.

Jitendra Upadhyay

Education & Activities

Udacity
Nano Degree in Data Science 2018

Indian School of Business Hyderabad
Master's in Business Analytics with A+ grade (Gold Medallist)
2019

Rajiv Gandhi Proudyogiki Vishwavidyalaya
Bachelor of Technology (B Tech) Hons 2009

Certifications

CFA Institute 2018
Certified Financial Analyst Level 2

Certifications & Awards

- HDP Apache Certified Spark Developer
- The Data Scientist Tool kit Program (JOHNS HOPKINS University) Cloudera.
- Part of a Cloud throne Event Winning team Rank 1
- Hackathon Event Winning team Rank 1 (theme to create a AI based traffic management system extensively used Graph Theo)