

## **Dr. UDAY MENON**

43 Windswept Road, Holmdel, NJ 07733 | H: 732 217 3150 | C: 646 522 1617 | [umenon@acm.org](mailto:umenon@acm.org)  
<https://www.linkedin.com/in/menonuday/>

### **Executive Profile**

Big Data architect interested in helping organizations acquire competitive advantage by leveraging their operational data to improve decision making. Expertise spans the full spectrum of skills from ETL to Analytics. Individual contributor who can conceive technology solutions and lead a team to implement and deliver them.

Also have a great interest in leveraging more than two decades of industry experience in the classroom, teaching graduate level courses in Data Analytics, Data Science and the full spectrum of Big Data technologies. Excellent communicator who can convey the essence of complex concepts using simple but powerful design oriented language, drawing from real-world examples. Towards that end, have also been certified to conduct training classes for clients of Datastax, the only vendor for Cassandra. Classes range from Operations using the Datastax Enterprise platform to Data Modeling, and Analytics using Apache Spark and Cassandra.

His current interests include developing a sports analytics app using Deep Learning technologies that he hopes will revolutionize the way tennis is taught at the all levels of the game.

### **Skill Highlights**

- Big Data (Spark, Solr, Elasticsearch, Hadoop MapReduce, NoSQL, Cassandra, Hive, Pig, Sqoop, Flume, Hbase, Drill)
- Data Science (Machine Learning, Statistical Inference, Python and R)
- Messaging (Kafka, Confluent, Tibrv, Solace, JMS)
- Object Oriented Technologies (Scala, Java, Smalltalk)
- Distributed Systems and Enterprise Architecture
- Relational Databases (SQLServer, Sybase, Oracle)

### **Core Accomplishments**

- Teach graduate level courses in the joint MSBA program offered by Columbia Engineering and Business School, bringing over two decades of industry experience into the classroom
- Completed technology assessment of a real time inventory platform for largest brick and mortar U.S. retailer and made recommendations for changes/enhancements in their Kafka messaging platform to improve throughput and latency
- Architected (design through implementation) an Apache Spark based compute-intensive solution for a leading Wall Street bank's risk management platform that fully leveraged the technology's parallel processing capability to deliver significant latency gains. Led an 8-man project in this effort which replaced an existing application that was unstable and slower.
- Key member of a hedge fund client technical architecture team that envisioned a migration path for their distributed applications to leverage big data technologies and transform the way they deliver

information to end-users. Hands-on involvement in all phases of the project from conception through vendor and technology selection through implementation and delivery

- Development Manager for Product Reference Data Distribution platform that feeds firm-wide trading and settlement systems. Managed a global team that was responsible for the platform
- Key development team member of an Enterprise Network Management startup that delivered a fully functional system for automatic discovery of network devices, servers, and software stacks and built a model of the enterprise network which helped root cause analysis of faults

## **Professional Experience**

### **Fulltime Faculty in Industrial Engineering & Operations Research      Sep 2018 to Present Columbia University – New York, NY**

- Teach Data Analytics and Supply Chain courses in the MS in Business Analytics (MSBA) joint program offered by Columbia Engineering and the Business School.
- Teaching Data Analytics, Computational Models In Supply Chain, Foundations of Data Science, Analytics In Practice and Analytics In Python in Spring 2024
- Taught 3 sections of Data Analytics in Fall 2023
- Taught Data Analytics and Computational Models In Supply Chain in Spring 2023
- Taught 3 sections of Data Analytics in Fall 2022
- Taught Data Analytics and Analytics in Practice in Spring 2022
- Taught Analytics in Practice in Fall 2021
- Taught Data Analytics in IEOR and GSAPP in Spring, Summer & Fall 2021
- Taught one section each of Data Analytics in Summer, Fall A, Fall B 2020
- Taught one section of Data Analytics for Real Estate in GSAPP in Fall B 2020
- Taught three sections of Data Analytics in Fall 2019
- Taught two graduate level courses, Data Analytics and Foundations of Data Science in Spring 2019
- Taught one section of Data Analytics in Fall 2018

### **Consultant Datastax – Santa Clara, CA**

**Jan 2018 to Aug 2019**

- Completed a technology maturity assessment of a Realtime Inventory platform for the largest brick and mortar retailer in the US. The platform uses Apache Kafka for messaging, Apache Cassandra for fast access to inventory lookup, ElasticSearch for adhoc querying and Hadoop for longterm storage and analytics. Proposed architecture changes to increase throughput and reduce read latency.
- Reviewed current computational model for a leading global fleet management company and advised them on their analytics workload including machine learning library selection, python programming best practices in a distributed computing environment and monitoring of spark jobs.

### **Managing Director Doran Jones, Inc. - New York, NY**

**Jul 2015 to Dec 2017**

- Replaced Impala based solution with Apache Spark execution engine (Scala implementation) as the heart of a risk management platform running on an on-prem cluster.
- Evaluation and selection of Hadoop distribution vendor. Export of RDBMS data using Sqoop into Hive tables and setup of Oozie update jobs to sync with source systems on a periodic basis. Creation of an Hbase cache for a Portfolio Management application which speeded up retrieval time of json data used to render graphs and tables from seconds to milliseconds. This also served as a replacement for in-memory cache which was unstable. Python implementation of Spark jobs running on a 5 node cluster which fetched fund performance data from data warehouse and supporting databases, assembled resulting json and stored in Hbase keyed by URL of browser query
- Confluent/Kafka platform implementation to track changes in RDBMS tables/views/queries, publishing to Kafka topics that triggered computes in pykafka subscribers.

## **Vice President**

**Jun 2009 to Jun 2015**

### **Barclays Capital - New York City, NY**

- Java Deduplication implementation of Product Reference Data from two source systems in order to migrate downstream consumers to a single source system without requiring changes to their interfaces
- Smart update of changes to product reference data attributes disseminated to downstream real time consumers resulting in reduction of total tibrv message volume by a factor of 7
- Big data repository of Compliance-mandated 7-years' worth of historical Reference Data using Hbase, Solr and Elastic search which achieved the twin objectives of cost-effective storage and ability to query for changes to data in arbitrary time windows
- Led a team to implement storage and retrieval of Price Time Series History information using a Hadoop Hbase NoSQL database capable of storing 40TB of data on a 5 node cluster and REST API to support browser based methods for data ingestion and querying.

## **Senior Developer**

**Feb 2008 to Feb 2009**

### **GFI Group - New York City, NY**

- Member of the Credit Trading System development team which supports a real-time electronic trading platform for CDS and Bond trades in the inter-dealer space.
- Implemented a Price Sheet Query Service which supports browser-based queries and returns an ordered collection of Interests with their current prices, organized by Workspace containing them. The results were used as part of a Bloomberg data feed provided by GFI
- Implemented instrumentation of API server code using AspectJ. The instrumentation involved injecting java bytecode with code that would collect metrics data (timestamp and context) and report to a Metrics Collector Service

## **Associate Director**

**Nov 2006 to Nov 2007**

### **UBS - New York City, NY**

- Implemented database access layer using Hibernate including writing hibernate mappings for all Java classes, writing delegates for these classes which are used by the object serialization
- Designed and implemented an object "diff" tool that invokes queries against two data servers (the targets of comparison) for all instances of business objects spanning Instrument, Parameter and supporting classes that together comprise the risk management suite.

## **Consultant**

**Jun 2000 to Nov 2006**

**JP Morgan - New York City, NY**

- Senior member of Investor Work Station application development team. 3-tier architecture comprising a Web front-end written in JavaScript & JSP, middle layer written in Java & C++,
- Provided CORBA interface for front-end JavaScript and JSP classes that need to access data from C++ servers.
- Implemented Java Data Access Framework that provides transparent data access from persistence layer (Sybase)

**Senior Developer****Jan 1997 to Jun 2000****Avesta Technologies - New York City, NY**

- Key architect of a Distributed Enterprise Network Management Software Application written in Smalltalk
- Developed Model-based Reasoning approach that employed Event Correlation and Impact Analysis to propagate a root-cause event through the network.
- Wrote Event dispatcher module which handled event parsing and interpretation

**Consultant****Feb 1994 to Nov 1996****Morgan Stanley - New York City, NY**

- Supported a Smalltalk based Enterprise Network Management system used to automatically discover entire network topology of routers, hubs, subnets, servers and applications running on them.

**Associate****Oct 1991 to Jan 2004****Booz Allen & Hamilton - New York City, NY**

- As member of Information Systems practice, worked on a variety of IT Strategy and implementation assignments for Fortune 500 organizations.
- Member of a 4-man developer team that implemented the first Smalltalk application which tracked sales trading for a major investment bank

**Adjunct Professor****Sep 1990 to Aug 1991****New York University - New York City, NY**

- Taught Introduction to MIS class to undergraduate students at Stern School of Business

**Education**

- **Ph.D.** Decision Sciences University of Illinois - Urbana Champaign, IL, USA - **July 1992**
- **MBA** IT and Marketing Indian Institute of Management(IIM) - Calcutta, India - **May 1981**
- **B. Tech.** Mechanical Engineering Indian Institute of Technology (IIT) – Kanpur, India – **June 1977**

**Certifications**

- Datastax Certified Professional on Datastax Enterprise
- Datastax Certified Professional on Apache Cassandra
- Cloudera Certified Hadoop Developer