

Course Title	Big Data Analytics				
Course Code	WSS552				
Course Type	Specialization (Elective)				
Level	Master (2nd Cycle)				
Year / Semester	2 or 3 Semester				
Teacher's Name	Christos Markides, PhD				
ECTS	10	Lectures / week	3	Laboratories / week	0
Course Purpose	<p>The term Big Data was coined in the early 2000s, where sciences like genomics, and nuclear physics first experienced the data explosion. Companies and organizations in the private and public sectors are capturing trillions of bytes of information. This information is the fruit of the information society, as data is collected from customers, suppliers, and operations, as well as millions of networked sensors embedded in the physical world. These devices are sensing, creating and communicating data from mobile phones and transportation systems, to CCTV systems and sensor networks. With a smartphone in every pocket, or a tablet in every purse posting information, and updating multimedia content over social media thus fuelling the volume of information generated at an even faster rate. To this end, data is no longer regarded as static or stale, that has fulfilled its purpose once it was collected. Data has become the raw material for business, and a tool to produce new form of economic value, yield innovation, and create new services.</p> <p>The purpose of this course is to provide students with a holistic approach to Big Data, the data model for Big Data, and examine the nature and requirements of a Big Data components, as well as Big Data as a platform. The course will then introduce the Apache Hadoop architecture, the Hadoop file system for distributed data storage, and highlights of MapReduce, Spark, Pig, and Hive as a framework of Big Data technologies as well as other tools for distributed processing of large datasets, using simple programming models. The course will address the shift from traditional relational database design for cluster systems to large scale NoSQL distributed databases, and the current trends for NoSQL database, as well as the future of Big Data. The course is complimented with practical and real-world examples. Finally, the course will examine the Big Data Analytics Lifecycle and review basic data analytic methods using R, Python, and Tableau, as an integrated environment for data integration, analysis, calculation, and visualisation for obtaining insight.</p>				
Learning Outcomes	<p>Upon successful completion of this course, students should be able to:</p> <ul style="list-style-type: none"> • Assess the state of Big Data adoption across a number of industry sectors. • Describe the Hadoop Architecture, and use the Hadoop file system and components. • Define and describe the major characteristics of NoSQL databases. • Deploy, manage, and interact with NoSQL databases using interfaces, programming languages, and queries. 				

	<ul style="list-style-type: none"> Analyse structured, unstructured, and IoT data for obtaining insight. Evaluate schemas for different types of data stores. Apply key concepts of Data Analytics Lifecycle to tackle Big Data problems. Develop technical skills to analyse data for data exploration (ingest, store and secure data). 		
Prerequisites	None	Corequisites	None
Course Content	<p>This module consists of the following chapters:</p> <ul style="list-style-type: none"> Chapter 1 introduces the fundamentals of Big Data and the main concepts of Data Mining. Chapter 2 explains the architecture of the Hadoop Framework and foundation core components Chapter 3 examines the architecture and characteristics of NoSQL databases Chapters 4 and 5 introduces Big Data Clustering platforms as general-purpose frameworks for cluster computing for data science and applications Chapters 6 and 7 expand general-purpose frameworks with the applications of NoSQL database and disparate storage needs Chapter 8 introduces the theory and application of Data Analytics and the Data Analytics Lifecycle Chapters 9 and 10 discusses the theory and application of Data Analytics methods Chapter 11 and 12 discusses methods and develops skills for data pattern discovery, data exploration, analysis, visualisation, and communication 		
Teaching Methodology	<p>The methodology used to conduct the course is structured around lectures and laboratory exercises, so that students gain theoretical knowledge as well as practical skills. The taught part of the course is delivered to the students with the help of computer presentations. Lecture notes and presentations are available through the web (e-learning platform) for students to use in combination with the textbooks. Furthermore, theoretical principles are explained by means of specific examples and for solving specific problems using practical examples.</p> <p>Lectures are supplemented with supervised and unsupervised computer laboratories. Laboratories will include demonstrations of taught concepts and experimentation with related technologies. Additionally, during laboratory sessions, students apply their gained knowledge and identify the principles taught in the lecture sessions by means of working on different tasks and problems. Students are also allocated exercises for homework, assignments and/or group project to improve both their individual skills and team work.</p> <p>In addition, a number of case study readings are also considered, in order to demonstrate the adoption, application, and practice of Big Data Analytics in various organizations and the industry. The case studies illustrate the characteristics of Big Data, as well as how Analytics help organizations to obtain insight.</p> <p>The course also includes a group project, in which students are expected to identify and work on specific Big Data Analytics, disseminate load, perform</p>		

	literature review in order to identify how similar problems were addressed elsewhere, design and implement a solution. Group projects are evaluated based on the project report and presentation.
Bibliography	<p>The following textbooks are associated with topics considered at various points throughout this course.</p> <p>EMC Education Services (2015), <i>Data Science and Big Data Analytics: Discovering, Analyzing, Visualizing and Presenting Data</i>, John Wiley & Sons, 1st ed.</p> <p>Krishnan K. (2013), <i>Data Warehousing in the Age of Big Data</i>, Morgan Kaufmann, 1st ed.</p> <p>Ladley J. (2019), <i>Data Governance: How to Design, Deploy, and Sustain an Effective Data Governance Program</i>, Academic Press; 2nd edition.</p> <p>White T. (2015), <i>Hadoop: The Definitive Guide</i>, O'Reilly Media, 4th ed.</p> <p>Sadalage P., Fowler M. (2012), <i>NoSQL Distilled: A Brief Guide to the Emerging World of Polyglot Persistence</i>, Addison-Wesley Professional, 1st ed.</p> <p>Lars G. (2011), <i>HBase: The Definitive Guide</i>, O'Reilly Media, 1st ed.</p> <p>Capriolo E., Wampler D., Rutherglen J. (2012), <i>Programming Hive</i>, O'Reilly Media, 1st ed.</p> <p>Anderson J. C., Lehnardt J., Slater N. (2010), <i>CouchDB: The Definitive Guide</i>, O'Reilly Media. Available [Online]: http://guide.couchdb.org/</p> <p>Grolemund G., Wickham H., (2016), <i>R for Data Science: Import, Tidy, Transform, Visualize, and Model Data</i>, O'Reilly Media, 1st Edition.</p> <p>McKinney W. (2017), <i>Python for Data Analysis: Data Wrangling with Pandas, NumPy, and IPython</i>, O'Reilly Media, 2nd Edition.</p> <p>Meier M., Baldwin D., and Strachnyi K., (2021), <i>Mastering Tableau 2021: Implement advanced business intelligence techniques and analytics with Tableau</i>, Packt Publishing, 3rd Edition.</p> <p>An extensive reading list of relevant academic papers:</p> <p>Sanjay G., Howard G., and Shun-Tak L., <i>The Google File System</i>, (SOSP '03). ACM, New York, USA, October 2003.</p> <p>Jeffrey D. and Sanjay G., <i>MapReduce: Simplified Data Processing on Large Clusters</i>, Commun. ACM 51, Jan. 2008.</p> <p>Chang F., Dean J., Ghemawat S., Hsieh W. C., Wallach D. A., Burrows M., Chandra T., Fikes A., and Gruber R., <i>Bigtable: A Distributed Storage System for Structured Data</i>. ACM Trans. Comput. Syst. 26, 2, Article 4, 2008.</p> <p>DeCandia G., Hastorun D., Jampani M., Kakulapati G., Lakshman A., Pilchin A., Sivasubramanian S., Voshall P., and Vogels W., <i>Dynamo: amazon's highly available key-value store</i>. In Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles (SOSP '07). ACM, New York, NY, USA, pp 205–220, 2007.</p>

	<p>Lakshman A. and Malik P., Cassandra: a decentralized structured storage system. SIGOPS Oper. Syst. Rev. 44, 2, pp 35–40, April, 2010.</p> <p>Massa S., and Testa S., Data warehouse-in-practice: exploring the function of expectations in organizational outcomes. Inf. Manage. 42, pp. 709-718, July, 2005.</p> <p>Wang J., and Liu B., Design of ETL Tool for Structured Data Based on Data Warehouse. In Proceedings of the 4th International Conference on Computer Science and Application Engineering 2020, ACM, New York, NY, USA, Article 119, pp. 1–5, 2020.</p> <p>Arora R., Pahwa P., and Bansal S., Alliance Rules for Data Warehouse Cleansing. In Proceedings of the 2009 International Conference on Signal Processing Systems (ICSPS '09). IEEE Computer Society, USA, pp. 743–747, 2009.</p> <p>Weber K., Otto B., and Österle H., One Size Does Not Fit All—A Contingency Approach to Data Governance. J. Data and Information Quality 1, 1, Article 4, 27 pages, June, 2009.</p> <p>Gopalkrishnan V., Steier D., Lewis H., and Guszczka J., Big data, big business: bridging the gap. In Proceedings of the 1st International Workshop on Big Data, Streams and Heterogeneous Source Mining: Algorithms, Systems, Programming Models and Applications (BigMine '12). ACM, New York, NY, USA, pp. 7-11, 2012.</p> <p>Sinaeepourfard A., Garcia J., Masip-Bruin X., and Marín-Torder E., Towards a comprehensive data lifecycle model for big data environments. In Proceedings of the 3rd IEEE/ACM International Conference on Big Data Computing, Applications and Technologies (BDCAT '16). ACM, New York, NY, USA, pp. 100–106, 2016.</p> <p>Tsai, CW., Lai, CF., Chao, HC. et al. Big data analytics: a survey. Journal of Big Data 2, 21, 2015.</p>
Assessment	<p>The Students are assessed via continuous assessment throughout the duration of the Semester, which forms the Coursework grade and the final written exam. The coursework and the final exam grades are weighted 40% and 60%, respectively, and compose the final grade of the course.</p> <p>Various approaches are used for the continuous assessment of the students, such as class participation and in class exercises, assignments, projects and tests. The assessment weight, date and time of each type of continuous assessment is being set at the beginning of the semester via the course outline. An indicative weighted continuous assessment of the course is shown below:</p> <ul style="list-style-type: none"> • 4 Online Quizzes (4 x 5% = 20 %) • 2 online participation activities (10%) • 1 marked assignment (group) project (10%) • 1 online / video presentation (10 %) • One final written exam (50%)

	<p>Students are prepared for final exam, by revision on the matter taught, problem solving and concept testing and are also trained to be able to deal with time constraints and revision timetable.</p> <p>The criteria considered for the assessment of each type of the continuous assessment and the final exam of the course are: (i) the comprehension of the fundamental concepts and theory of each topic, (ii) the application of the theory in solving related problems and (iii) the ability to apply the above knowledge in complex real-life problems.</p> <p>The final assessment of the students is formative and summative and is assured to comply with the subject's expected learning outcomes and the quality of the course.</p>
Language	English