Teaching IoT (Internet of Things) Analytics

Jai W. Kang
Rochester Institute of Technology
152 Lomb Memorial Drive
Rochester, NY 14623
585-475-5362
jai.kang@rit.edu

Qi Yu
Rochester Institute of Technology
152 Lomb Memorial Drive
Rochester, NY 14623
585-475-6929
qi.yu@rit.edu

Erik Golen
Rochester Institute of Technology
152 Lomb Memorial Drive
Rochester, NY 14623
585-475-4409
efgics@rit.edu

ABSTRACT
The rapid proliferation of Internet of Things devices around the world has led to a major increase in demand from industry for students equipped with the skills necessary to make continued advances in this area. Consequently, advanced analytical skills are in urgent need to capitalize the massive amount raw data collected by various IoT devices. To address the market demand in IoT and Analytics, the Information Sciences and Technologies department at the Rochester Institute of Technology has proposed an advanced certificate in IoT Analytics that extends across its three Master of Science degree programs. The central focus of this work is a presentation of this advanced certificate program that is designed to encompass the four pillars of IoT, namely Sensors, Communications, Computing Devices, and Analytics.

Categories and Subject Descriptors
K.3.2 [Computers and Education]: Computer and Information Science Education – curriculum, information systems education.

Keywords
Information sciences and technologies; big data; cloud computing; curriculum; data analytics; internet of things; network.

1. INTRODUCTION
The Information Sciences and Technologies (IST) department at Rochester Institute of Technology (RIT) strives for keeping the degree programs up to date with technology advances to prepare students for their job hunting and future professional career. We adopted one of the disruptive technologies, *Cloud Computing,* as a database course in 2009 [11] and reported its progress in 2011 [10]. In 2013, we revised our master's program as an *Analytic* centric program [14], and its success story appeared in [15]. While recommending *Security* requirements to be embedded in the courses across two master's degree programs in IST and NSA (Networking and Systems Administrations), we also considered the modern computing landscape of three key building blocks: *Internet of Things (IoT), Cloud Computing,* and *Big Data* [16]. This paper emphasized that students need to understand security requirements not only within each of the three building blocks but the interactions thereof.

SIGITE'17, October 4–7, 2017, Rochester, NY, USA

ACM ISBN 978-1-4503-5100-3/17/10…$15.00
https://doi.org/10.1145/3125659.3125689

Gartner anticipates that there will be 26 billion Internet connected things by 2020 resulting in the generation of massively amounts of data that need to be stored and processed [18]. As IoT technology advances, the value can be realized from analyzing the data generated by the connected things. In other words, a real insight is generated by the intersection of IoT and Analytics, which is referred as *IoT Analytics* [17]. [7] refers IoT as a prominent driver to the fourth Industrial Revolution (IR) that will have impacts throughout the business and industry continuum around the world. *IoT Analytics* motivates our next curricular update in the IST department at RIT.

As part of pursuing details of IoT Analytic course topics, a number of IT related recent curricula research articles reveal many IoT course offerings: Teaching IoT Concepts [1]; IoT Design of Experiment [4]; Open Source Platform Design [5]; Course Design for IoT Using Lab of Things from Microsoft Research that includes Data Mining for future plan [13]; Active Learning in Open Elective Courses [19]; Raspberry Pi: An Effective Vehicle in Teaching the IoT [26]; Integrating IoT into STEM Undergraduate Education [9]; and Teaching the IoT [6] which is the only graduate course in these findings. None of these courses associate teaching to the interaction of IoT and Data Analytics as IoT Analytics. A few courses include their teaching approaches by Active Learning, Project-Based, and Problem & Project Based Learning.

As far as IoT Analytic course topics are concerned, we start defining the pillars of IoT with their technologies. The four pillars are Sensors, Communications, Computing Devices, and Analytics. Together, these pillars constitute the progression of IoT data from its rawest form to knowledge discovered from collected data. While raw sensor data progresses to a gateway and then to the cloud, it is important to recognize which locations we are able to apply the analytics to generate real insight: *Edge, Data Stream, and Data Lake*.

To ensure success, it is also important to determine the correct format for a curricular structure. In this paper, we provide an overview of the IST department, which justifies why we are at a unique position to teach IoT analytic related topics. We then identify a number of curricular structures to teach these topics, including a new master program, a concentration within an existing master program, and an advanced certificate focusing on IoT analytics. We discuss the major benefits of an advanced certificate to show why it provides the best option for us.

The remainder of the paper is organized as follows. We begin by introducing the four pillars of IoT and describe the basic concepts and functionalities of each in Section 2. Section 3 describes our curricular choice of teaching IoT Analytics. Section 4 discusses the curriculum of an advanced certificate consisting of four courses, which will be offered in a 2+2 format with the first two as core courses and the last two as elective courses. Section 5 discusses how the new program will be assessed before concluding the paper in Section 6.

2. Pillars of IoT

Our Advanced Certificate in IoT Analytics is structured around what we have deemed to be the pillars of IoT, which include *sensors, communications, computing devices, and analytics* as shown in Figure 1. Together, these pillars constitute the progression of IoT data from its rawest form to knowledge discovered from collected data.

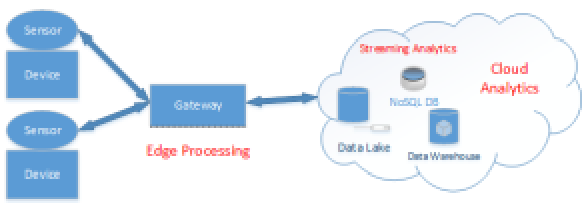

Figure 1. IoT Pillars

2.1 Sensors

With respect to IoT, sensors are miniature devices that monitor a particular phenomenon of interest [8]. For example, in a Smart Home, some typical phenomena include temperature, humidity, and motion sensing for monitoring the environment and keeping track of its inhabitants. Personalized healthcare monitoring devices like the FitBit extend beyond a static environment and accompany their users by sensing motion through accelerometers, gyroscopes, and magnetometers, depending upon their complexity. Regardless of the application, sensors convert electrical input signals to units called *counts* that are related to the sensitivity of the sensor. Using the known sensitivity of the sensor, the counts are then scaled to provide human understandable units such as degrees Celsius.

Sensor data must be periodically transmitted from the source to a sink node, or *aggregator* [2]. The responsibility of the aggregator is to collect data from various IoT devices and direct the data towards remote processing resources in the cloud and eventually, back to the IoT devices. As such, communications play an important role in moving data from the sensor that generated it to the processing resource that will perform analytics. To limit locality restrictions on IoT devices, wireless communication between IoT devices and an aggregator is standard.

2.2 Communications

In most IoT applications, the distance between an IoT device and an aggregator is short, which allows for single hop wireless communication standards such as Bluetooth Low Energy, Zigbee, and 802.11 in either infrastructure or ad-hoc mode to be used. In terms of curriculum, the amount of prerequisite networking knowledge required for students is minimal since the networking is accomplished on a point to point basis with little complexity [2]. For rare cases where large areas must be sensed, such as monitoring moisture across farmland extended tens of acres, multi-hop communication may be needed, at which point more advanced protocols are required. Regardless of the distance between IoT devices and an aggregator, once sensor data has reached aggregator, it will reach its processing resources over a wired connection on the Internet.

2.3 Computing Devices

Computing devices consist of three main types in IoT. These include data collection platforms, aggregators, and remote processing resources [8]. Data collection platforms comprise one or more sensors, resource limited processor boards, low capacity memory modules, and wireless communication devices. These platforms are constructed in a variety of ways from prototype boards, such as an Arduino or Raspberry Pi to custom micro controllers to single board computers. Common among all of these approaches is that data collection platforms have relatively finite processing resources and memory, which means that they must possess the capability to offload sensor data with sufficient frequency to support the application at hand.

As previously stated, data aggregators act as gateways for IoT device data to reach remote computing devices and for the remote computing devices to command IoT devices after processing data. Data aggregators are therefore equipped with both wireless and wired communication devices and have much higher processing capabilities and onboard memory than data collection platforms. When IoT data reaches the aggregator, it must identify which device sent the data and then determine which remote processing resources it should be routed to and vice versa. Depending upon the technological maturity of the data collection platform, an aggregator may be as simple as a Linksys access point to as complex as a customized data collector housed in a microcomputer.

Once the aggregator has directed IoT device data to the appropriate remote processing resource in the cloud, analytics can begin. Remote processing resources may range from high performance computing clusters to servers equipped with GPUs to virtualized processing containers.

The progression of raw sensor data reaching an aggregator and then its ultimate destination of the remote processing resource leads up to the final IoT pillar, analytics. Within the analytics pillar, we consider three types - *Edge Processing*, *Streaming Analytics*, and *Cloud Analytics* [12].

2.4 Analytics

While the raw sensor data make progresses to reach an aggregator and then to its ultimate destination of the remote processing resource, it is important to recognize at which locations we are to apply the analytics in order to generate real insight from the IoT ecosystem. There are three potential locations: *Edge, Data Stream, and Data Lake*. The types of analytics depend on factors like volumes of datasets, data types, quality of data, privacy concerns, real-time or batch processing needs.

Edge Processing: When the IoT sources generate large datasets, it increases latency and may be infeasible to transfer due to poor or underpowered network connections. A solution to these types of issues is to process and analyze data directly at or near the source – also known as *Edge Processing* [12]. For example, modern jet engines generate up to 1 TB of data per flight. Even though traditional aggregators are equipped with limited computing power, both hardware gateway vendors like Cisco and Intel, and server vendors like Dell offer smart gateways with more computing power including more storage and even analytic capabilities. The large dataset can be processed in a smart Gateway to send only average values or outliers to the cloud.

Alternatively, the aggregated data can be stored locally in the smart Gateway and transferred to the cloud when the network connection is better [3].

Streaming Analytics: Forrester defines *Streaming Analytics* platform as a "software that can filter, aggregate, enrich, and analyze a high throughput of data from multiple disparate live data sources and in any data format to identify simple and complex patterns to visualize business in real-time, detect urgent situations, and automate immediate actions [23]." When data streams generated from the sensors reach the remote processing resource, the Cloud, they can be applied to either a real-time system or a batch system based on system requirements. The real-time system analyzes the streaming data to perform analytics without storing them on a data store, which is called *Streaming Analytics*. This streaming analytic includes performing statistical analysis of data in motion or applying a predictive model to find any deviations in the real-time stream data. The predictive models are built offline using machine learning algorithms based on the historical training data.

Cloud Analytics: When the sensor data reach the Cloud for batch processing, they arrive in a Data Lake, which is driven by the Hadoop platform and conceptually similar to a Data Mart or Data Warehouse. But data in a Data Lake holds its raw format until it is needed for *Cloud Analytics* or other needs like OLAP (Online Analytical Processing) cube analysis. Appling ETL (Extract, Transform and Load) processes to data in the lake, they can be updated in either a data warehouse for the OLAP queries or a NoSQL database for Natural Language Processing (NLP) as well. Data analytics techniques including statistical analysis, machine learning and visualization can be applied to the IoT datasets in the Data Lake [12].

3. ADVANCED CERTIFICATE in IoT ANALYTICS

In this section, we start by providing an overview of the IST department, which helps justify why we are at a unique position to teach IoT analytics related topics. We will then identify a number of curricular structures to teach these topics, including building a new degree program, forming a concentration within a program, and creating a new advanced certificate focusing on IoT analytics. We will discuss the major benefits of the advanced certificate to show why it provides the best option for us.

The IST department is the home of three undergraduate and three graduate programs. Since the IoT and analytics related courses are at the graduate level, we will briefly introduce the three graduate programs of the department. In particular, the M.S. in Information Sciences and Technologies (IST) puts a strong emphasis on data science with a focus on how data is analyzed, managed, and visualized in the modern computing industry. The M.S in Networking and System Administration (NSA) covers both classical theories/concepts on wired and wireless network modeling and analysis as well as emerging computing and networking technologies, such as cloud computing and Internet of Things. Finally, the M.S. Human-Computer Interaction (HCI) explores the design, evaluation, and implementation of interactive computing systems. The ability to offer these three graduate programs demonstrates that the faculty of the department has the required expertise to cover all the important components in an IoT system that includes the underlying (wired and/or wireless) network infrastructure, data acquisition techniques and UI design patterns to collect data from IoT devices, and analytical tools/models for knowledge discovery from the collected data.

The desired technical expertise from the faculty coupled with the success of running three graduate programs puts the IST department in a unique and strong position to offer IoT analytics. To ensure success, it is also important to determine a right format for a curricular structure. There are three possible choices: a separate degree program, a concentration with a program, and an advanced certificate. A master level graduate program at RIT consists of 30 credit hours, which is roughly equivalent to 10 courses (most graduate courses in RIT are worth of 3 credits). Even though we can reuse some of the courses from department's existing graduate programs, a decent number of new courses may still need to be developed to make a complete degree program. This will be a rather time consuming process. Even after these courses are developed, new faculty members may need to be recruited to teach these courses as the department has already fully utilized its personnel to cover the existing courses. Getting additional faculty lines demands a lengthy administrative approval process. A much more efficient way is to offer IoT analytics related topics to group a number of relevant courses and make them a concentration. A concentration typically consists of 2-4 elective courses to provide enough depth of study in a given program. However, there are two major limitations with a concentration. First, it is only available to the students in one graduate program. Due to the interdisciplinary nature of the IoT analytics, it is hard to decide which of the three existing programs is the best host for this new concentration. Second, while a concentration offers a way to organize students' plan of study, it is not explicitly reflected in their transcripts or diplomas. However, students may want this important skillset to be formally recognized, which may benefit their job hunting and future professional career.

Finally, an advanced certificate is comprised of a four-course sequence that allows a student to develop expertise in a particular field. An advanced certificate offers a number of key benefits, making it our best choice to offer IoT analytics. First, it only consists of four courses, making both course development and delivery much easier than a 10-course degree program. Second, it is standalone and independent from other degree programs. But students in other degree programs (e.g., IST, NSA, and HCI) can pursue the advanced certificate simultaneously. Third, the certificate provides a formal way to recognize the skills and special training that a student receives in this field. Last, besides serving the current students, the advanced certificate also provides an effective means to attract new students into RIT and the IST department. Some students may start by joining the advanced certificate and then extend their study to a full M.S. program, such as IST, NSA, and HCI based on their interest.

4. CURRICULUM

The curriculum of the advanced certificate consists of four courses, which will be offered in a 2+2 format with the first two as the core courses and the last two as the elective courses. The two core courses aim to build the foundation in IoT analytics, which all the students are required to take. The first core course will be a newly developed IoT foundation course that focuses IoT side of the advanced certificate, covering the fundamental concepts in IoT technology. The second core course will be an existing data analytics course being offered by our MS-IST program, titled "Analytical Thinking (ISTE-600)". This course will focus on the analytics side of the certificate, covering fundamental data mining techniques, including various kinds of supervised and unsupervised data mining models, data preprocessing, and model evaluation.

The elective courses will build upon the fundamental skills that students develop through taking the core courses and further extend them through advanced training. We are currently considering the following five elective courses including two advanced analytics courses: Data-Driven Knowledge Discovery (ISTE-780) and Visual Analytics (ISTE-782), and two IoT related courses: Data Acquisition and Analysis in IoT (NSSA-yyy), Prototyping Wearable & IoT Devices (HCIN-720). Four of these courses are exiting courses currently being offered by one of our three graduate programs. Reusing existing courses not only minimizes the effort for course development and delivery, but also facilitates and motivates students in our current graduate programs to take the advanced certificate. Table 1 summarizes topics covered in the certificate by core and elective courses followed by course descriptions.

Table 1. IoT Topics vs. Courses

	Core Courses		Electives (Choose 2)			
	NSSA-xxx	ISTE-600	ISTE-780	ISTE-782	NSSA-yyy	HCIN-720
Hosting Degree Programs	NSA (new)	IST	IST	IST	NSA (new)	HCI
IoT Pillars/Topics						
1) Sensors:						
Programming sensors	x					x
API	x					x
Hardware	x					x
2) Communications:						
Wireless Communications	x				x	x
Networking Protocols & Standards	x				x	x
Info Privacy & Security	x					
3) Computing Devices:						
Data Collection Devices	x				x	x
Aggregator	x				x	
Remote Processing Resources	x				x	
4) Analytics:						
Data Mining (Supervised/ Unsupervised)		x	x			
Data Collection		x	x	x	x	
Data Preprocessing		x	x	x	x	
Statistical Analysis		x	x	x	x	
Big Data (Hadoop, NoSQL)			x			
Cloud Computing	x		x			
Visualization		x		x	x	x

4.1 Core Courses

4.1.1 NSSA-xxx Foundation of IoT

This course provides students with an overview of the area of IoT and hands-on laboratory experiences to put their knowledge into practice. The overall IoT architecture is first presented with the information flow from data collection devices to aggregators to remote processing resources and back again. Throughout the remainder of the course, the underlying technologies for these architectural building blocks are discussed in detail, with laboratory exercises given on both the Raspberry Pi and Arduino Uno as exemplary platforms, culminating in a final project where students develop a prototype IoT system of their own. To motivate the data analytics side of IoT, example applications are shown throughout the semester that highlight Edge Processing, Streaming Analytics, and Cloud Analytics, and how these processing techniques and resultant feedback are integrated into the IoT architecture.

4.1.2 ISTE-600 Analytical Thinking

This core course focuses on the analytics side of the IoT Analytics certificate, covering fundamental data mining techniques, including various kinds of supervised and unsupervised data mining models, data preprocessing, model evaluation and visualization.

The Critical Thinking Community [12] defines all thinking by eight elements: "Whenever we think for a purpose within a point of view based on assumptions leading to implications and consequences. We use concepts, ideas and theories to interpret idea, facts, and experiences in order to answer questions, solve problems, and resolve issues. Thinking, then 1) generates Purposes, 2) raises Questions, 3) uses Information, 4) utilizes Concepts, 5) makes Inferences, 6) makes Assumptions, 7) generates Implications and 8) embodies a Point of View."

Students customize the above analytical thinking approaches to solve data mining problems following the Cross Industry Standard Process for Data Mining [20]: 1) business understanding, 2) data understanding, 3) data preparation, 4) modeling, 5) evaluation, and 6) deployment.

Students work in teams on a problem of their choosing that is interesting, significant and relevant to applying data mining algorithms and techniques to real-world problems like IoT. Students use the Weka [25], other data mining software, and Tableau [21] for visual analysis and presentations. Student teams construct excellent stories with interactive and dynamic dashboards for their projects with the help of these visual analytic tools.

Topics covered in Fundamental Analytics: Fundamental data mining techniques, including supervised (e.g., Decision Tree, Rule-based, Nearest-Neighbor, Naïve Bayes, Artificial Neural Network & Ensemble Models) and unsupervised (e.g., Association, Cluster Analyses & Anomaly Detection) data mining models; data preprocessing; model evaluation; and visualization.

4.2 Elective Courses

4.2.1 ISTE-780 Data Driven Knowledge Discovery

ISTE-780 is one of the advanced elective courses in the analytics domain of the MS-IST program. This course provides advanced training to students in data analytics, with a focus on statistical learning approaches in the context of the data-driven knowledge discovery process. The main objectives of this class are to

1. Model and understand complex datasets using statistical

learning tools that discover useful information and knowledge from large-scale datasets by conducting both supervised and unsupervised learning.
2. Scale statistical learning algorithms with powerful, distributed, and cloud-based systems (e.g., Apache Hadoop and Mahout) to handle large-scale datasets.
3. Learn data analytics languages (R and Python) and apply statistical packages (R and scikit-learn) to tackle real-world data analytics problems

Main topics of the course include both state-of-the art supervised and unsupervised statistical learning models. On the supervised learning side, it covers regression models, such as best subset regression, ridge regression, LASSO, and principle components regression, and classification models, such as logistic regression, linear discriminant analysis, support vector machine, and random forest. On the unsupervised learning side, it covers key clustering models, such as k-means and spectral clustering, and dimensionality reduction techniques, such as principle component analysis and latent dirichlet allocation.

4.2.2 ISTE-782 Visual Analytics
The main thrust of visual analytics is to discover patterns in data that were previously obscured prior to visualizing it [24]. Latent patterns are rarely seen when data is confined to a spreadsheet or database. From an IoT perspective, visual analytics is important in that collections of disparate sensors are likely to be analyzed and the ability for humans to make sense of multiple inputs is necessary. For example, in a smart home, a user would like to gain a complete understanding of their monitored environment whether or not they are home to allow them to make necessary adjustments to temperature, humidity, and other phenomena. Furthermore, sensor readings are space and time varying, as these are the key features in Geographical Information Systems (GIS) and remote sensing applications. This space-time visualization is also critical in visualizing IoT data since these devices may span longer distances than a single home and consist of long time spans. This course covers space-time variations in data, interaction with visualizations, human cognition of visualizations, and extends to application areas of GIS, IoT, and cyber security and how visualizations inform practitioners in those areas.

4.2.3 NSSA-yyy Data Acquisition and Analysis in IoT
To bridge the gap from IoT data generating platforms to data analytics, this course provides students with the opportunity to learn how to generate data sets of their own from existing IoT devices in order to produce meaningful results from analytics. This is in stark contrast to typical student experiences, where the student will locate a data set online that has been analyzed countless times and may not contain the data necessary to produce meaningful results. The course begins with a discussion of practical issues associated with gathering data from sensors, include sensitivity of sensing devices, power consumption, data storage and in the case of multi-sensor systems, time synchronization between devices and differing sampling rates. Fundamentals of data acquisition are then discussed with classical design of experiments presented, including hypothesis testing and data sampling, so that students understand the importance of collecting the “correct” data so that features can be constructed. During the feature construction process, students are shown that data acquisition and performing analytics is an iterative process where the analytics often informs the experimenter about whether or not the features are sufficient in number and in content.

4.2.4 HCIN-720 Prototyping Wearable and IoT Devices
This course focuses on rapidly prototyping and evaluating the utility of IoT devices, wearable or otherwise, specific. For the purposes of this course, the student will recognize an IoT device as the proliferation of hardware, software, and resultant data generated by the device. As this course is offered out of the Human Computer Interaction MS degree program, its focus is on the user experience with the device being prototyped. This includes the physical interaction between user and prototype, such as how understandable a smart thermostat may be to the average person. Prototyping skills learned in this course range from 3D printing, laser cutting, sewing, and modeling to capacitive sensing, actuation, and electronics theory. To bring the prototypes to life, students learn event-driven programming skills required in programming an Arduino or Photon and external communication through Bluetooth Low Energy or WiFi. Data is represented in the course through visualization and web technologies for display, such as the REST web API and node.js.

5. ASSESSMENT
In this section, we identify a few program goals and student learning outcomes (see Table 2), which will be used to assess the Advanced Certificate in IoT Analytics. Assessment data will be collected and reviewed every three years. The student learning outcomes are assessed from student activities in coursework against a rubric designed for that outcome. The benchmark is that 80% of the students will achieve competence in that outcome. Two of those are assessed using coursework in the core and one from a selected elective course.

Table 2. Program Assessment

Program Goals	Student Learning Outcomes	Data Source / Measurement
1. Design IoT systems to collect, communicate, and aggregate sensor data.	Create a prototype IoT system that consists of all the major computing devices and collect data for simple analysis.	NSSA-xxx *Foundations of IoT:* Rubric used to assess the fundamental knowledge in IoT
2. Apply specialized analytical and technical skills for IoT data analysis	Preprocess and mine sensor data using different types of IoT analytics techniques	ISTE-600 *Analytical Thinking:* Rubric used to measure effectiveness of IoT data analysis
3. Design IoT information services to enhance the value of raw IoT data of various types	Demonstrate advanced data acquisition and/or analytical skills in the IoT domain	An advanced project in a student-selected elective course - Project is assessed via rubric

Besides assessing student-learning outcomes, we also plan to evaluate the attractiveness of the proposed advanced certificate by keeping track of the number of capstones that our current MS students choose to work in IoT analytics. The capstone is the culminating experience that a student must complete in order to obtain a MS degree. Choosing a capstone topic in IoT analytics demonstrates that a student has not only gained sufficient technical expertise but also developed strong interest in studying and working in this area.

6. CONCLUSION

To meet industry demand for skilled practitioners in IoT and data analytics, the IST Department at RIT has crafted a new IoT Analytics advanced certificate program that extends across its three MS programs. Students who earn this certificate are expected to carve out a valuable niche in that their skills will run the gamut of the four pillars of IoT, namely Sensors, Communications, Computing Devices, and Analytics. Through its two core courses, students will acquire the background knowledge in IoT architecture and analytics needed to delve further into their chosen interest areas that are covered by the four elective courses that span the three MS programs. Future work includes additional electives, such as an advanced networking and computing technologies course for implementing cloud analytics applications.

7. REFERENCES

[1] Ali, Farha. 2015. Teaching the Internet of Things Concepts. In *Proceedings of the WESE'15: Workshop on Embedded and Cyber-Physical Systems Education* (WESE'15), Martin Edin Grimheden (Ed.). ACM, New York, NY, USA, , Article 10 , 6 pages. DOI: http://dx.doi.org/10.1145/2832920.2832930

[2] Atzori, L., Iera, A., & Morabito, G. 2010. The Internet of Things: A Survey. *Computer Networks*, 54(15), 2787-2805.

[3] Baars, H., & Ereth, J. 2016. From Data Warehouses to Analytical Atoms–The Internet of Things as a Centrifugal Force in Business Intelligence and Analytics.

[4] Chen, G. W. 2016. The Design of Experimental Course for Internet of Things Based on ZigBee. In *Intelligent Human-Machine Systems and Cybernetics (IHMSC), 2016 8th International Conference on* (Vol. 2, pp. 341-344). IEEE.

[5] Dobrilovic, D., & Zeljko, S. 2016. Design of open-source platform for introducing Internet of Things in university curricula. In *Applied Computational Intelligence and Informatics (SACI), 2016 IEEE 11th International Symposium on* (pp. 273-276). IEEE.

[6] Förster, A., Dede, J., Könsgen, A., Udugama, A., & Zaman, I. (2017). Teaching the Internet of Things. *GetMobile: Mobile Computing and Communications*, *20*(3), 24-28.

[7] Geng, H. (Ed.). 2017. *Internet of Things and Data Analytics Handbook*. John Wiley & Sons.

[8] Gubbi, J., Buyya, R., Marusic, S., & Palaniswami, M. 2013. Internet of Things (IoT): A Vision, Architectural Elements, and Future Directions. *Future generation computer systems*, 29(7), 1645-1660.

[9] He, J., Lo, D. C. T., Xie, Y., & Lartigue, J. 2016. Integrating Internet of Things (IoT) into STEM undergraduate education: Case study of a modern technology infused courseware for embedded system course. In *Frontiers in Education Conference (FIE), 2016 IEEE* (pp. 1-9). IEEE.

[10] Holden, E.P., J.W. Kang, Anderson, G.R., D.P. Bills, M. 2011. Databases in the Cloud: A status report. ACM SIGITE Conference 2011, October 20-22, 2011 West Point, NY.

[11] Holden, E.P., J.W. Kang, D.P. Bills, M. Ilyassov. 2009. Databases in the Cloud: A work in progress. ACM SIGITE Conference 2009: pp 138-143, October 22-24, 2009 Fairfax, VA.

[12] Jaokar, A. 2017, Data Science for Internet of Things: Strategies for Product Development Using Streaming, Sensor Fusion and Deep Learning.ajit,jaokar@futuretext.com

[13] Jeong, G. M., Truong, P. H., Lee, T. Y., Choi, J. W., & Lee, M. 2016. Course design for Internet of Things using Lab of Things of Microsoft Research. In *Frontiers in Education Conference (FIE), 2016 IEEE* (pp. 1-6). IEEE.

[14] Kang, J. W., Holden, E. P., & Yu, Q. 2014. Design of an analytic centric MS degree in information sciences and technologies. In *Proceedings of the 15th Annual Conference on Information technology education* (pp. 147-152). ACM.

[15] Kang, J. W., Holden, E. P., & Yu, Q. 2015. Pillars of Analytics Applied in MS Degree in Information Sciences and Technologies. In *Proceedings of the 16th Annual Conference on Information Technology Education* (pp. 83-88). ACM.

[16] Kang, J. W., Yu, Q., Holden, E. P., & Oh, T. H. 2016. Security Requirements Embedded in MS Programs in Information Sciences and Technologies. In *Proceedings of the 17th Annual Conference on Information Technology Education* (pp. 77-82). ACM.

[17] Markkanen, A., & Shey, D. 2014. The intersection of analytics and the internet of things. *IEEE Internet of Things*.

[18] Modi, C., Patel, D., Borisaniya, B., Patel, A., and Rajarajan, M. 2013. A survey on security issues and solutions at different layers of Cloud computing. *The Journal of Supercomputing*, 63(2): 561-592.

[19] Raikar, M. M., Desai, P., & Naragund, J. G. 2016. Active learning explored in Open elective course: Internet of Things (IoT). In *Technology for Education (T4E), 2016 IEEE Eighth International Conference on* (pp. 15-18). IEEE.

[20] Shearer, C. 2000. The CRISP-DM model: the new blueprint for data mining. *Journal of data warehousing*, *5*(4), 13-22.

[21] Tableau: http://www.tableausoftware.com/.

[22] The Critical Community: Elements and Standards Learning Tool. Retrieved May 12, 2014, from http://www.criticalthinking.org/pages/elements-and-standards-learning-tool/783.

[23] The Forrester Wave: Big Data Streaming Analytics, Q1 2016. Retrieved June 1, 2017 from http://sqlstream.com/wp-content/uploads/2016/12/Forrester-Streaming-Analytics-Report-2016.pdf

[24] Thomas, J. J. 2005. Illuminating the Path: The research and development agenda for visual analytics. *IEEE Computer Society*.

[25] Weka: http://www.cs.waikato.ac.nz/ml/weka/.

[26] Zhong, X., & Liang, Y. 2016. Raspberry Pi: An Effective Vehicle in Teaching the Internet of Things in Computer Science and Engineering. *Electronics*, *5*(3), 56.