Predictive Data Analytic Approaches for Characterizing Design Behaviors in Design-Build-Fly Aerospace and Aeronautical Capstone Design Courses

Dr. Krishna Madhavan, Purdue University - West Lafayette

Dr. Krishna Madhavan is an Assistant Professor in the School of Engineering Education at Purdue University. He is Co-PI and Education Director of the NSF-funded Network for Computational Nanotechnology (nanoHUB.org which serves over 330,000 global researchers and learners annually). Dr. Madhavan was the Chair of the IEEE/ACM Supercomputing Education Program 2006. In January 2008, he was awarded the US National Science Foundation (NSF) CAREER award for work on learner-centric, adaptive cyber-tools and cyber-environments. He was one of 49 faculty members selected as the nation's top engineering educators and researchers by the US National Academy of Engineering to the Frontiers in Engineering Education symposium. Dr. Madhavan leads a major NSF funded effort called Deep Insights Anytime, Anywhere (DIA2) that attempts to characterize the impact of NSF and other federal investments in the area of science, technology, engineering, and mathematics education using interactive knowledge mining and visual analytics for non-experts in data mining. DIA2 is currently deployed inside the NSF and is already starting to affect federal funding policy. Dr. Madhavan also served as Visiting Research Scientist at Microsoft Research, Internet Services Research Group. His research has been published in Nature Nanotechnology, IEEE Transactions on Computer Graphics and Applications, IEEE Transactions on Learning Technologies, and several other top peer-reviewed venues. Dr. Madhavan currently serves as PI or Co-PI on federal and industry funded projects totaling over $20M.

Dr. Michael Richey, The Boeing Company

Dr. Michael Richey is a Boeing Associate Technical Fellow currently assigned to support technology and innovation research at The Boeing Company. Michael is responsible for leading a team conducting engineering education research projects that focus on improving the learning experience for students, incumbent engineers and technicians. His research encompasses, Sociotechnical Systems, Learning Curves, and Engineering Education Research. Additional responsibilities include providing business leadership for engineering technical and professional educational programs. This includes development of engineering programs (Certificates and Masters) in advanced aircraft construction, composites structures, systems engineering, product lifecycle management and digital manufacturing. The educational programs and research focus on practical understanding of human learning and the design of technology-enhanced learning environments and promoting global excellence in engineering and learning technology to develop future generations of entrepreneurially-minded engineers. This is achieved by partnering and investing in educational initiatives and programs between industry and institutions of higher learning. Under Michael's leadership, The Boeing Company has won the multiple Awards for Excellence and Innovation for their industry academic partnerships and joint programs Michael has served on various advisory groups including, the editorial board of the Journal of Engineering Education, Boeing Higher Education Integration Board, American Society for Engineering Education Project Board and the National Science Foundation I-UCRC Industry University Collaborative Research Center Advisory Board. Michael has authored or co-authored over 25 publications in leading journals including Science Magazine, The Journal of Engineering Education and INCOSE addressing topics in large scale system integration, learning sciences and systems engineering. Michael often represents Boeing internationally and domestically as a speaker - presenter and has authored multiple patents on Computer-Aided Design and Computer-Aided Manufacturing, with multiple disclosures focused on system engineering and elegant design.Michael holds a B.A and M.Sc. from ESC Lille in Program Project Management and Ph.D. from SKEMA Business School with a focus on Engineering Education Research.

Mr. Barry McPherson, The Boeing Company

Barry McPherson is the Senior Business Partner for education programs in the Technical and Professional Learning Solutions organization within The Boeing Company. His focus is on building customized learning solutions for both on-hours and off-hours programs that bring Boeing quality and innovation into the

learning space. Barry manages a multimillion dollar research and delivery portfolio focused on advanced manufacturing (e.g. composites) and Product Lifecycle Management (PLM) disciplines, providing technical excellence to Boeing engineers regarding technical foundations and cutting-edge applications.

Predictive data analytic approaches for characterizing design behaviors in design-build-fly aerospace and aeronautical capstone design courses

Abstract:
Predictive data models and interactive visualizations can be highly effective in understanding workload and skills assignment issues within design-build-fly teams in the aerospace industry. Capturing data that is needed to build predictive models in usable forms and then subsequently applying appropriate data mining techniques to derive insights from such data is a significant challenge. The ultimate goal of our work is to understand design behaviors among engineers that can lead to cost reductions and expediting product development in extremely complex engineering environments. The present study, pioneered by a large US aerospace company working with educators at 5 major engineering schools in the US, engineering education researchers, and practicing engineers, is a first step towards achieving this overall vision. In this paper, we characterize how engineering students enrolled in a senior capstone course interact and perform on complex engineering tasks commonly seen in the aerospace industry. We describe our instrumentation methodology and the data architecture for an associated analytics platform. We use course clickstreams, social networking and collaborations as the basis for our observations.

1. Introduction

"Design is the first step in manufacturing, and it is where most of the important decisions are made that affect the final cost [and indeed performance] of a product"[1] (p. 505). The entire aerospace and aeronautics industry relies heavily on the ability of engineering teams to design, test, and execute not only individual components, but also more importantly on their ability to translate these components to function effectively within entire systems. "Digital technologies are enabling a direct correlation between what can be designed and what can be built, thus bringing to the forefront the issue of the significance of *information*, i.e. the issues of production, communication, application and the control of information […]"[2] (p. vi).

Not only is the design of extremely complex systems like aircrafts driven to a great extent by digital technologies, but also more importantly these technologies are increasingly used to support distributed engineering teams that function in socially cohesive structures. Madhavan and Lindsay (2014)[3] point out that information technology has ushered in the realization that engineering problem solving is an increasingly socially-driven and supported enterprise. They state that "for example, the design of the Boeing 777 [aircraft] by engineers who were distributed in multiple parts of the world using modeling and simulations as a primary problem-solving modality is a watershed transformation in engineering practice" (p. 635)[3]. Information technology is fundamentally changing engineering practice in the workplace[4,5,6]. More importantly, the instrumentation available within these digital systems is capable of providing insights about how professional engineers design, work, manufacture, and support extremely complex systems. But most times, such performance and behavioral data derived from industrial systems is proprietary by nature.

"Engineering organizations are increasingly under pressure to perform more efficiently with fewer people. To manage this, organizations need to understand what skills, knowledge and behaviors they need from engineers who have to practice in a global information society [minor spelling changes to conform with US norms]" (p. 145)[7]. Understanding how engineers within a company design, work, and solve complex problems provides tremendous competitive advantage to those possessing those insights. Thus, most major engineering companies are extremely reticent to openly publishing such data about the design and execution behaviors of their engineering teams. In the absence of these insights, it is extremely difficult to fully characterize the behaviors that lead to better product design, increased productivity, and decreased product costs. Thus even if there is high quality instrumentation within the digital environments that support industrial design and manufacturing, it is difficult to develop characterizations of engineering behaviors as manifested in the environmental factors that digital environments can now capture. While there are numerous studies that shed light into the nature of engineering practice in the workplace[8, 27], there are very few studies, if any, that can provide a full and thorough analysis of engineering practice as manifested within the digital design environments where the work actually occurs.

1.1. Focus of this paper

The present study, pioneered by a large US aerospace company working closely with educators at 5 major engineering schools in the US, engineering education researchers, and practicing engineers, is a first step towards characterizing engineering problem solving behaviors in the aerospace industry. In this paper, we characterize how engineering students enrolled in a senior capstone course interact and perform on complex engineering tasks commonly seen in the aerospace industry. We see the senior capstone course as a microcosm of the engineering behaviors we expect to see in the real world. There is significant literature to show that well designed capstone courses offer a realistic engineering context for fostering practice relevant design behaviors[9,10,11]. The core purpose of this paper is to describe our instrumentation methodology and the data architecture for an associated analytics platform. Further, we show the various types of analytic artifacts that would be needed in describing the engineering behaviors of students fully. We intend to transfer these analytic capabilities to training and education activities with this large US aerospace company with the express goal of driving on-the-job training costs down and increasing productivity. In the next section, we provide the context of the study and a description of the capstone design course.

2. Context of this study

The ultimate goal of our work is to understand design behaviors among engineers that can lead to cost reductions and expediting product development in extremely complex engineering environments such as the aerospace industry. The present study describes the first step taken to address this overall vision. Going beyond this goal of fully characterizing problem-solving behaviors among professional engineers as manifested in digital environments, our research vision is to build predictive data models and interactive visualizations that can be highly effective in understanding workload and skills assignment issues within design-build-fly teams in the aerospace industry. Capturing data that is needed to build predictive models in usable forms and then subsequently applying appropriate data mining techniques to derive insights from such data is a significant challenge.

The above stated goal poses a very high barrier of entry to solving this problem. One of the few viable pathways to addressing the above challenge is to establish a test bed that is defined, designed, and administered by an industry-led team, wherein learners work hand-in-hand with industry experts, academic researchers, and data scientists to elicit the type of design behaviors that reflect real world engineering practice in the aerospace industry. This allows us to develop, test, and refine the instrumentation methodology, data architectures, analytics, and visualization approaches before interfering with the day-to-day work within an organization. In the context of our work, a program called AerosPACE was developed not only as a senior capstone course, but also to serve as a test bed.

AerosPACE is an engineering education program developed by a large US aerospace company. The primary goal of this program is to bridge the gap between theory and application, (and to help students understand the process of designing, building, and flying an unmanned aerial vehicle (UAV) capable of assisting first responders. As students engaged in design activities, a second goal was to develop an instrumentation methodology and data architecture needed to fully characterize industry relevant engineering design behaviors as manifested in the digital environments. Multi-disciplinary, multi-university teams consisting of students from 5 major US universities participated in a two semesters, year-long capstone project. These courses have been effectively offered starting 2013. The third cohort of student teams is now experiencing this capstone course. This enables us to gather a significant amount of data related to design behaviors that form the basis for many of our data architecture decisions.

To illustrate this, the highly distributed nature of the design process and collaboration between students at different universities are major themes of the project. It is expected that each design will address technical areas of aerodynamics, materials, propulsion, manufacturing, structures, and controls among others. Major milestones include a Mission Concept Review, Preliminary Design Review, Critical Design Review/Production Readiness Review, Flight Readiness Review, and Post Launch Assessment Review in addition to a flight demonstration. The overall theme of the UAV's mission is to help various first responders protecting citizens in this country and across the world. First responders fulfill various missions, many of which can benefit from the use of small-unmanned aerial vehicles. As a part of this project, students define specific missions they will design their vehicle for to support first responders.

Since AerosPACE would also serve as a test bed for developing predictive models to address design behaviors in the aerospace industry, it needed to exhibit some core characteristics in order to establish the viability of our methods, data architectures, and analytic approaches. These characteristics were agreed upon based on numerous expert conversations and discussions. They are described below briefly:

> ***Flattening organizational learning curves:*** One of the critical features of AerosPACE* must be that - from a business perspective, it allows researchers in understanding the aggregate components and implications of organizational learning curves (ability of an

entire organization to learn and gain experience about a product or process)[12, 13, 14], including the impact of knowledge transfer on product performance. The flattening of the organizational curve is in effect the true intended effect. Flatter organizational learning curves allow companies to reach profitability faster and also increase overall product quality and performance. Therefore, the methodology we implement must allow us to understand pathways through which knowledge dissemination can be increased effectively.

Learning mimics true engineering practice in the aerospace industry: The test bed must allow engineering educators to take on the significant challenge of incorporating real world engineering competencies into their core curriculum. Engineering students have to increasingly be trained to work in extremely distributed teams that solve problems as a cohesive unit using state-of-the-art digital tools. Therefore, as much as feasible, the capstone course needed to utilize tools and services currently used in the aerospace industry for training and organizational learning needs.

Enable cross-pollination of industrial and academic practices: AerosPACE must allow the adoption of industrial practices by supporting mechanisms for incorporating industrial innovations into its core learning components. Additionally, it must enable the flow of educational innovations back into industrial practice.

In the next section, we describe the instrumentation designed to observe various aspects of this capstone course. Further, we also provide an overview of the data architecture that forms the basis of our analytics platform.

3. Methods – Instrumentation and Data Architecture for AerosPACE

The challenge for this study begins with instrumenting the design environment effectively. When engaging in the build-design-fly engineering process, students typically have to interact with a number of online learning and design environments (for example, learning management system, design environments for the aircrafts, simulation environments to test the design, and specification documentation capture systems). Developing an architecture needed to address this data pipeline is the first aspect that this paper addresses in significant depth. Figure 1 provides a simplified overview of the architecture.

Just as in the engineering workplace, students enrolled in the AerosPACE capstone design course interact with a wide range of online learning and design environments. Step 1 (top left box of Figure 1) shows all students initially beginning their interactions with some unified learning environment. Although Sharepoint[15] is used commonly to launch capstone design projects in the aerospace industry, this initial launch point could be any range of environments. Once the student is introduced to the course, they routinely interact with course content (step 2) that resides in a learning management system such as Blackboard[16] or even a MOOC platform (in the context of AerosPACE we use edX[17] as a content platform). Since the course contains a significant number of complex design components, students on a regular basis use synchronous and asynchronous communications to interact with each other (step 3). Throughout the design process, students maintain multiple design documents that track their design specifications and build requirements (step 4).

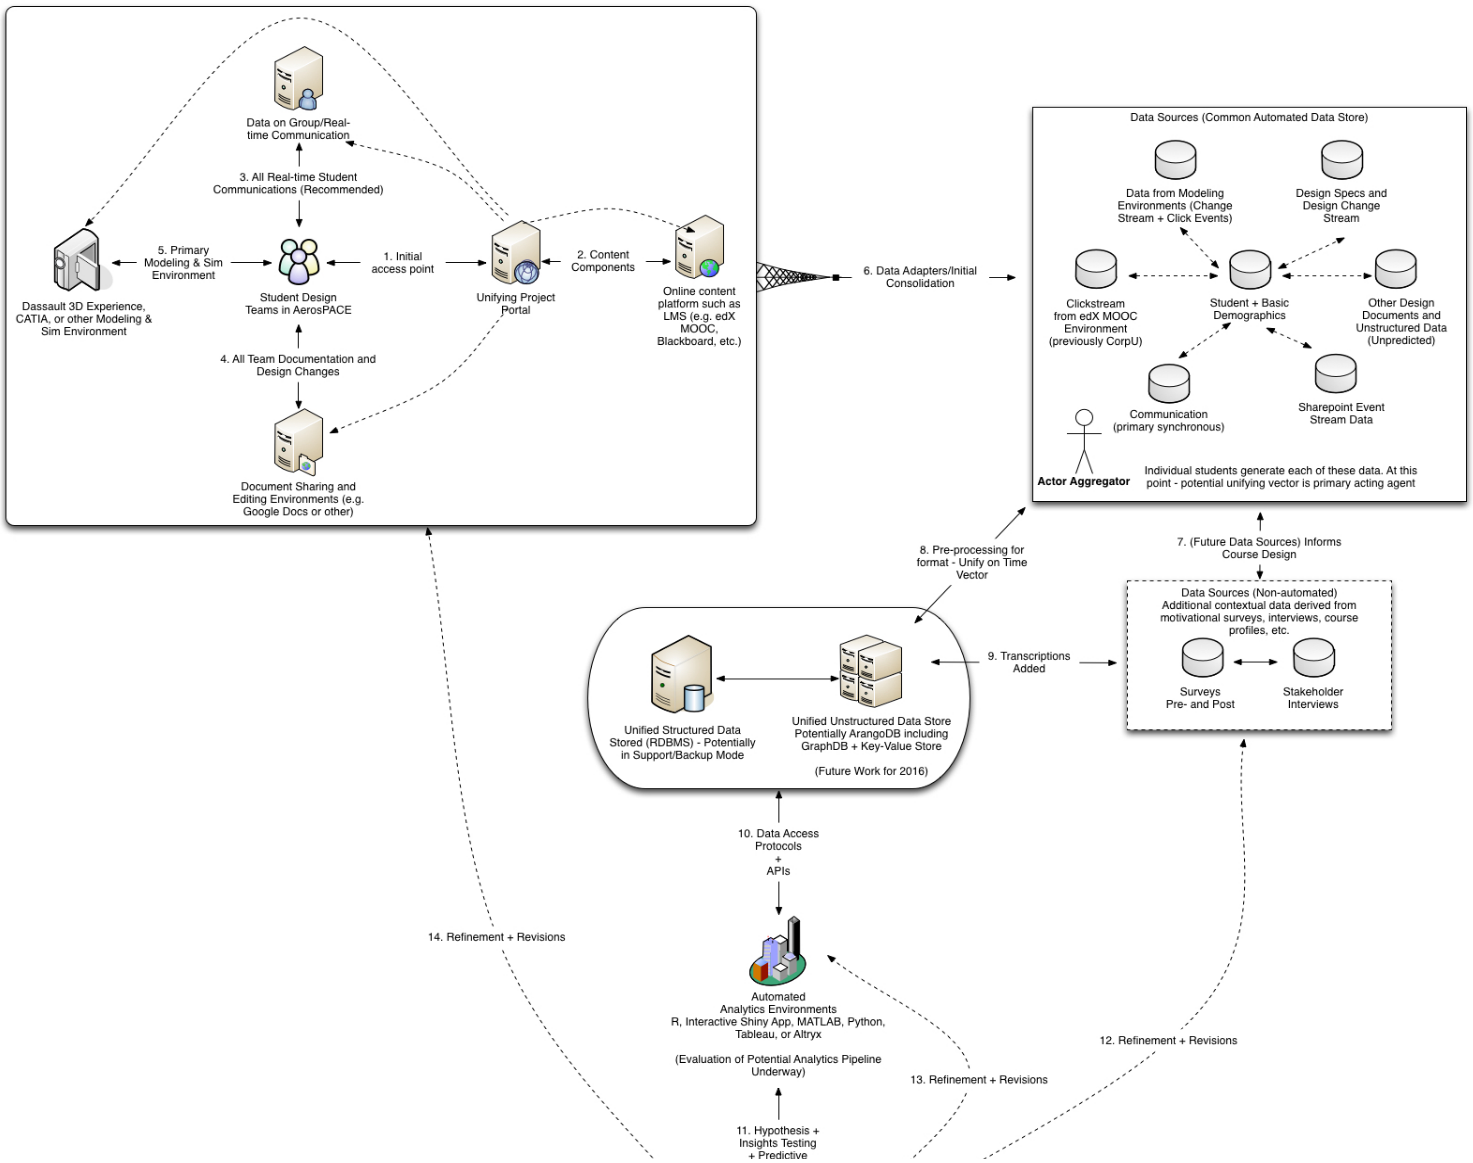

Figure 1. Simplified data and analytics platform architecture for AerosPACE

One of the core design activities that students engage in throughout the project is the use of 3D modeling and simulation environments (step 5) to develop *in silico* prototypes of the aircrafts they eventually will build and test. Essentially steps 1 – 5 show the various types of data and environments that students routinely experience during a design, build, test capstone course in the aerospace industry. These environments are very similar to the types of environments that professional engineers use within the aerospace company. This is one of the primary reasons why the course interactions are setup as shown here.

Students consistently generating a wide range of data types including, 3D models, design documents, social interaction data, course content interactions, and performance data. From an analytics and predictive modeling perspective, it is incredibly difficult to utilize these data in performance models or any type of decision models without transforming them into some unified schema. In step 6, we apply a variety of data filters to first sort the data into specific data types and then transform all these data to focus on individual learners (or actors in the context of our data architecture). At this point – our system is being designed to automatically parse all of the individual data and store them into a pre-processed form. The top right box in Figure 1 shows the automation. Once this step is complete, our system is designed to allow the addition of data points that provide better context for the learner data (step 7) that we already have pre-processed. Such contextual data could range from motivations behind taking the course, course profiles, and other pre- and post-data related to course performance.

The AerosPACE test bed collects data on student and faculty interactions from a variety of sources. This data represents, in a broad sense, records of actions taken in multiple web-based environments, communications shared between individuals, and surveys administered to students. We use clickstream data as a core part of our analytics effort. Madhavan and Richey (2016) define clickstream "[…] as a proxy for user behavior that not only records that a click was made by the user on an electronic object, but also captures the associated metadata, such as the system time when the click was made (for standardization), the object on which the click event occurred, the dwell time on the clicked entity, and a few other navigational details. Insofar as clickstreams embody behaviors, they are also a manifestation of learner intent" (p. 8)[18].

With data coming from several different streams, information must be blended together carefully to capture a complete picture of the student design behavior evolution and development over time. Secondly, these clickstream data also contain a mapping over time of students' interactions with faculty and industrial partners and time series distribution of skills and collaborative messages – essentially data that are ideal for social network analyses[19]. Additionally, text mining and web log mining techniques allows researchers to gain deep insights on the major discussion topics. Step 8 shows a key organizational phase required by any analytic systems that may be layered on top of the datasets – namely, unifying the data around EPOCH time or UNIX time. This essentially enables time-series based analyses. At this point, we also add transcripts and any appropriate information about the surveys to the data (step 9). Our data is composed of broadly two kinds of data – structured and unstructured data[20]. Therefore, our system utilizes a database type that can maintain structured, semi-structured, and unstructured data types. Furthermore, we are transitioning to a database system (e.g. ArangoDB[21]) that can maintain graphing relationships also.

At this point in the workflow, we are in the process of writing a series of data access protocols and application programming interfaces that use JSON to interact with our visual analytics layer (indicated in step 10). The visual analytics layer is critical to our system as this is where many of our insights are generated and is discussed in greater detail in the next section. From steps 11 through 14, various researchers interact with our system to derive insights, build predictive models, and provide direct feedback to both instructional designers and students who utilize the system.

4. Visual Exploration of Learner Behaviors and Implications for Predictive Modeling

As show in Figure 1, visual exploration of the data in step 10 is a critical stage. Easy exploration of learner data emanating from electronic environments as the learners are engaged in a design activity has never been easy. In order to enable visual exploration of all design behaviors, we are experimenting with multiple methodologies, tools, and platforms for visualizing learner data. Our goal is to eventually transition to a close-to-real-time visual analytic capability. In this section, we provide details based on the various platforms we have been testing.

4.1. Tableau

Tableau[22] is a versatile tool that can link to various different files on different types of servers, but in order to create meaningful visualizations, the data usually needs to be manipulated in some third-party program before it can be used in Tableau. For some of the visualizations in this project, the raw data were manipulated in Excel and brought into Tableau through a SQL Join within Tableau, which required no extra coding. Tableau is capable of handling relatively large chunks of data. Furthermore, the visualizations can be easily customized with just a few clicks, the visualization can be customized with appropriate colors, markers, and legends. However, it is critical to note that using Tableau as a visualization platform requires use of other data analytic environments to augment Tableau's capabilities. Tableau is a visual presentation platform, but not necessarily visual analytic platform aimed completely at non-experts in data mining. These visualizations allowed us to derive insights from time-series data on user interactions with the edX platform where all AerosPACE learning content is based.

Figure 2 shows digital objects accessed by users over time, organized by months with a single action. The time resolution can be adjusted into years, months, week number, weeks and days with a single action. The graph shows anonymized student IDs sorted on number of times a student accesses objects – and further in descending order and course sections broken down by Month vs. ID. Figure 2 shows that students systematically revisit topics when interacting with learning materials related to aerospace. These topics that students revisit on a regular basis could hold the key to recognizing threshold topics that are critical for students to master certain design factors. Figure 3 is a modified form of figure 2 that shows how intensely students use certain objects. It also shows that students repeatedly visit certain learning materials with varying intensity.

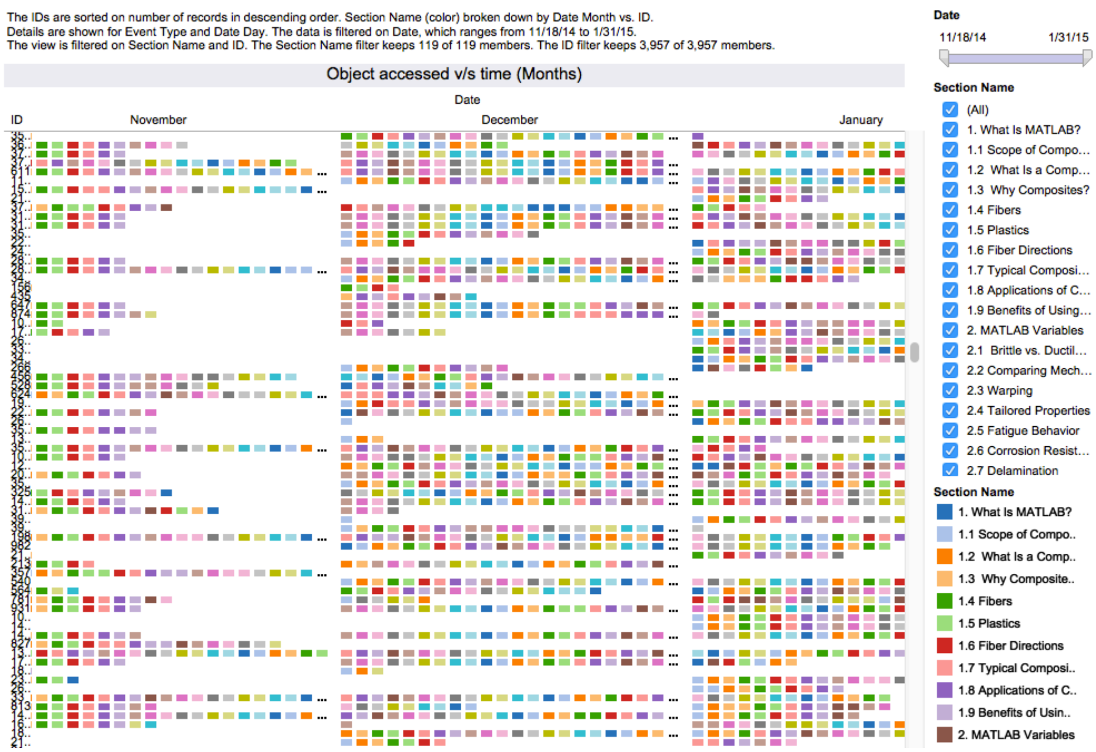

Figure 2. Visual patterns of users' interactions with digital objects in a learning environment

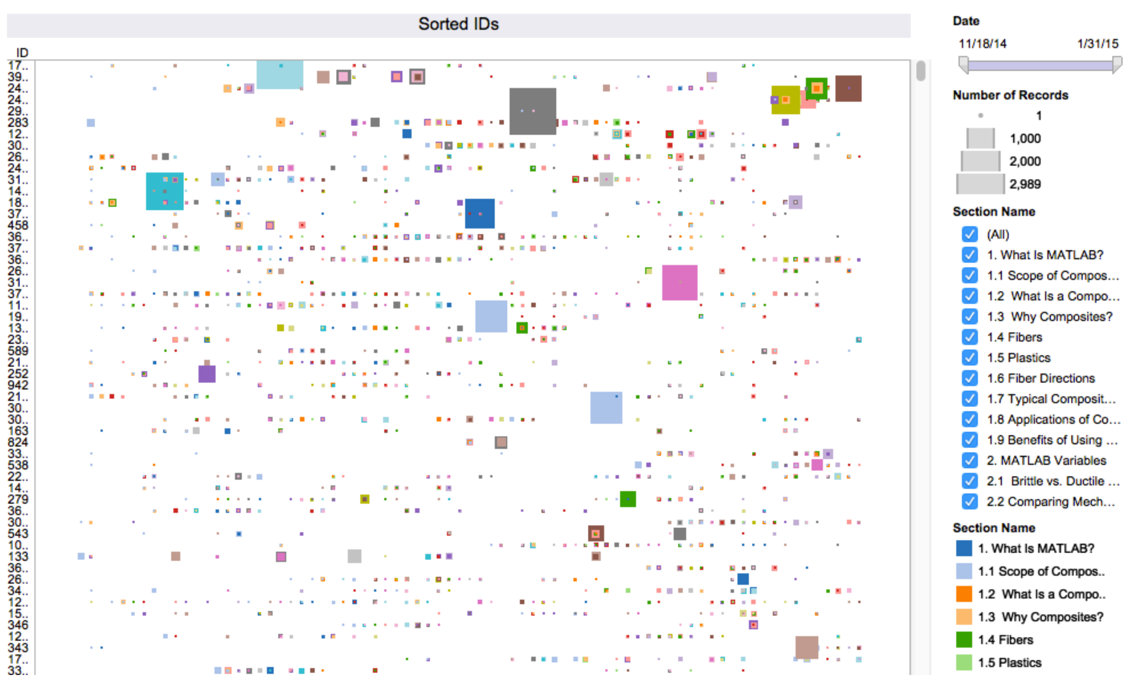

Figure 3. Intensity of learning content use and revisiting course content over time

The power of simple visual exploration cannot be underestimated. For example, figure 4 below is a simple line graph that shows that students view content for longer periods of time initially, then as the course progresses, over time students change their behavior and settle at a constant – but lower rate of content viewing. The spikes observed in figure 4 are a critical indication that at certain points in the semester, depending on the design task that students are engaged in, the intensity of content viewing increases in a spike.

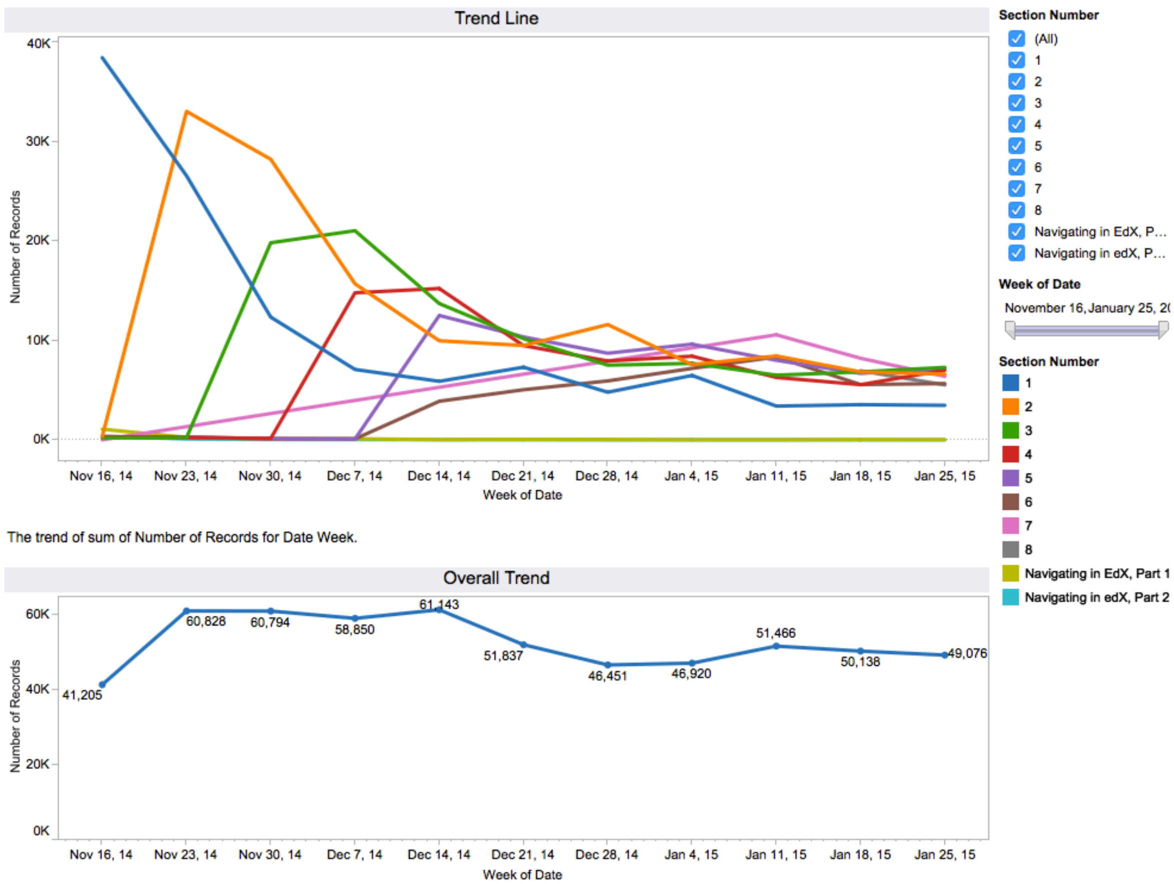

Figure 4. Content viewing trends over time showing intense initial use and then settling to a lower constant rate.

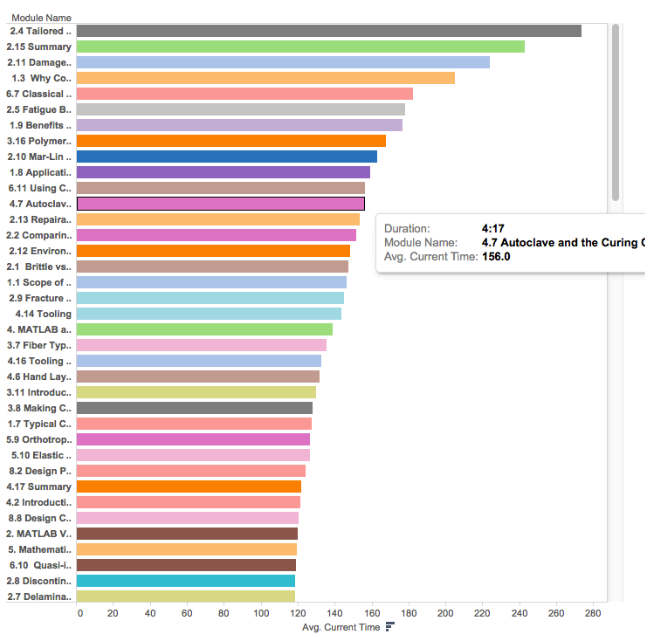

Figure 5 (Left). Pause behavior of students when viewing course content.

In figure 5, we demonstrate that visual exploration can also show content that may require students to invest more time to understand. The graph shows, in descending order, content components on which students paused the most. This graph shows a proxy for content difficulty and maybe quite useful in structuring materials in level-appropriate ways. For example, here our visual explorer shows that content component 2.4 on Tailored Properties was perhaps particularly troublesome for students as it registered the most pause events. The point here is that using visual exploration, we are now able to shed completely new light into how learning content behaves vis-à-vis learner behavior and cognitive demands.

While these types of visual explorations are critical, they are by far not sophisticated enough to develop deeper insights needed to build predictive models. In order to accomplish more in-depth and more comprehensive analyses of behaviors that will eventually feed a predictive model, we need to resort to more powerful analytic tools such as the R-Studio/ShinyApps platform. We discuss this next.

4.2. R-Studio/ShinyApp

R is an open-source programming language popular in the data science community. R-Studio[23] is an integrated development environment for R that is also free and open-source. R presents many advantages, being a free and open-source software with no license restrictions. It offers robust cross-platform compatibility and can be used on GNU/Linux, Mac OS X, and Microsoft Windows environments in both 32- and 64-bit configurations. Over 4,800 library packages are currently available for use, including many powerful packages for matrix manipulation, statistics, and graphical representations. As such, R is the most comprehensive software in this group for statistical analysis and visualization of results. R can natively accept input from a variety of sources, including those relevant to the AerosPACE project.

Because R is a command-line environment, it requires more advanced computer programming skills than the alternatives, which means there can be a steep learning curve. Analysis in R requires a lot of front-loaded effort to establish a workflow for data that is either unstructured, inconsistently structured, or poorly structured. Many critical analytic components of this study are implemented in a fully automated framework under R giving access to the analysis via a web application. This application allows researchers to interact with the results permitting executives and decision makers to go deeper into the training data. Our work also lowers significantly the information management barriers in how engineers are trained to participate in production-oriented teams.

4.2.1. Social interaction data analytics in AerosPACE: As engineers engage in the design process, they constantly interact with each other. It is important to point out that in design, build, test teams within the aerospace industry, each team member plays a critical role because they bring complementary but varied skillsets needed to complete the design, construction, and testing of the aircraft. However, traditionally, most aircraft construction teams are staffed with all necessary skillsets needed for completing the projects from day 1. Although there is a phased design and construction process, the skills needed within the project are usually pre-assigned. In AerosPACE, we were able to incorporate survey data that asked participants to rate themselves on specific skill dimensions. Based on learners' individual rating we were able to classify learners as being focused on CAD, Manufacturing, Propulsion, Testing, Controls, and so on. We then used our analytics platforms to understand when specific skills were needed based on interaction patterns within a given team. Figure 6 shows that for each of the skill areas, the workloads vary over time. During some phases of the project, certain skill sets are completely not engaged in the design process on any team. This essentially shows that perhaps there is a better workload assignment model that could be used on design, build, test teams within the aerospace industry. We are continuing further investigations along this direction by collecting more validation data.

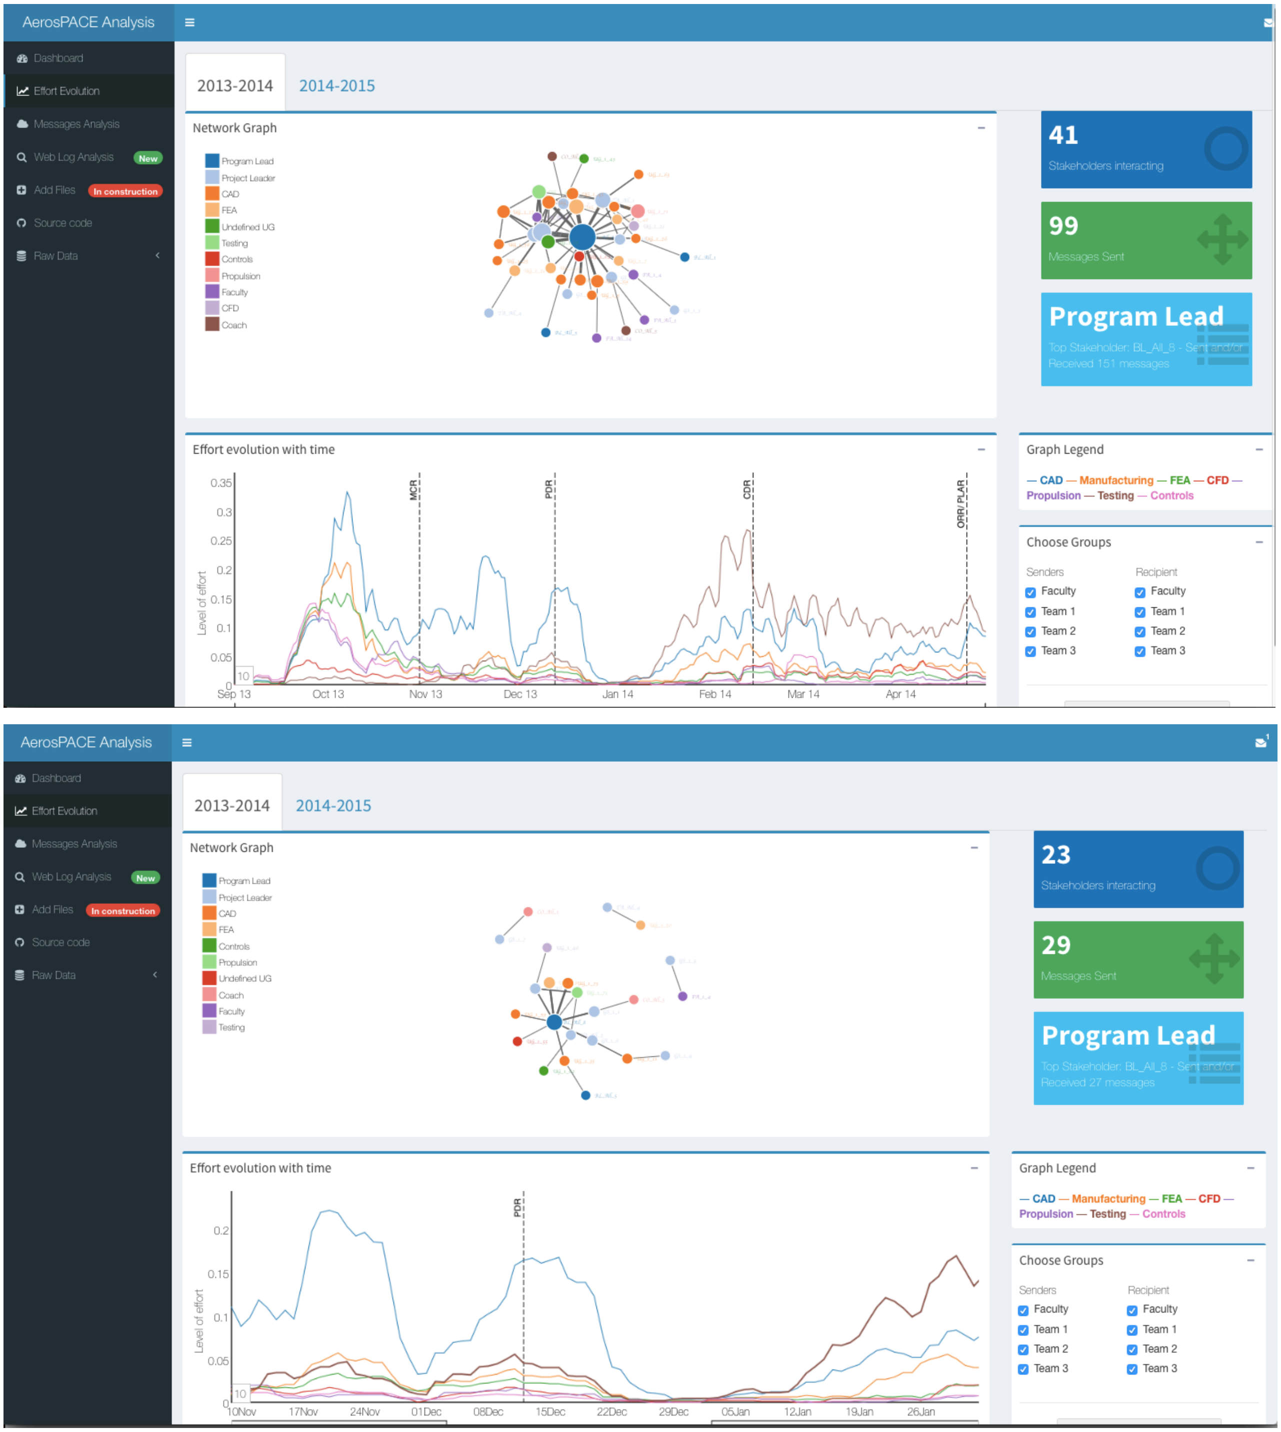

Figure 6a (Top) and 6b (Bottom). Workload distribution for specific skills over various time periods of the project. Social interactions show that not all skill sets are equally needed during all stages of the project. 6b (Bottom) shows a slightly constrained time period.

4.2.2. Clustering for learner behaviors: One of the key components that is needed to build a predictive model is the ability to group similar users together into appropriate clusters. The reason that clustering is critical is because it allows for personalization of learning and design pathways. It is widely known that personalization of learning is a National Academy of

Engineering grand challenge[24]. Identifying patterns of learner behaviors and classifying learners into specific clusters essentially allows us to build decision engines that can direct users based on the usage cluster to which they belong. To this end, the AerosPACE analytics platform has begun to include a two key clustering algorithms – K-Means Clustering[25] with the ability to dynamically adjust k-values and Hierarchical Clustering (Ward)[26]. Other additional clustering methods are actively being explored and implemented. The goal is to allow a team of data scientists to converge on a handful of clusters that are prototypical for engineers engaged in design, build, test activities in the aerospace industry. Figure 7 below showcases the clustering capabilities of our analytics platform. It clearly shows that engineering students perform a series of activities and they can be clustered based on their design behaviors. Furthermore, the bottom panel of Figure 7 shows that we can control the clustering based on specific time periods so that we can personalize design pathways in much more fine-grained ways using a combination of clustering and time series analyses.

Figure 7. Clustering capabilities within AerosPACE analytics platform featuring K-Means with dynamic k-value adjustment and Hierarchical Clustering (Ward).

In this section, we have attempted to provide a brief overview of how visual exploration of learner data can provide a new way to understand engineering design practice. We have discussed two very different approaches to data mining and have outlined the capabilities of each system. In the next section, we will briefly discuss predictive model construction.

4.3. Predictive Models for Understanding Design Behaviors in the Aerospace Industry

Thus far in this paper, we have framed the components needed in order to eventually build a set of predictive models that will eventually feed a decision engine. With all these components in place, we can now begin to identify a set of key factors that would form the core of the predictive model. Factor identification and prioritization is necessary to building strong predictive models.

At this point in the process, our workflow will split the data into a training dataset and a validation dataset. The training dataset will be used to systematically used for feature identification and for training a series of machine learning algorithms. A discussion of each of these machine-learning algorithms is outside the scope of this paper. However, we will also be evaluating these algorithms in ensemble mode, where the output of one algorithm will be used as input for the next algorithm. The idea is to be able to develop machine learning techniques that will be able to fully recognize (to the extent possible) typical design behaviors and their precursor behaviors. Once the models have been trained, we will then use the validation data set to ensure that robustness and accuracy of our models. The final stage of this work will be the utilization of these predictive models in decision-support mode to enable better efficiencies throughout the design, build, test, and fly process.

5. Conclusions

AerosPACE research is completely aligned to the corporate roles and responsibilities of the Technical Fellowship within the US aerospace company (blinded) through a) supporting the company's business strategies by ensuring technical excellence across the enterprise in our people, technologies, processes, tools, and products; b) Expanding the company's technical skills and performance by improving the acquisition, retention, knowledge, and utilization of our technical workforce for business success; and c) Participating in representing the company's technology interests to the outside world customers, general public, academia and government. This research partnership is investigating designs for technology-enhanced teaching, learning and assessment that connect opportunities for formal and informal learning and support an aeronautics workforce culture that is adaptive to change. Conceptually, AerosPACE and Engineering Education Research offer a completely cohesive approach to understanding and characterizing design behaviors within the aerospace industry.

Acknowledgments

This work was supported in part through a research contract from The Boeing Company to Purdue University. We also acknowledge the efforts of graduate students Charles Sese, Paul Branham, and Shweta Singh in building and testing the visual exploration components of this project.

Bibliography

1. Boothroyd G. Product design for manufacture and assembly. Computer-Aided Design. 1994;26(7):505-520.
2. Kolarevic B, ed Architecture in the Digital Age. New York, NY: Spon Press; 2003.
3. Madhavan K, Lindsay ED. Use of Information Technology in Engineering Education. In: Johri A, Olds BM, eds. Cambridge Handbook of Engineering Education Research. New York, NY: Cambridge University Press; 2014:633 – 654.
4. Black SE, Lynch LM. How to Compete: The Impact of Workplace Practices and Information Technology on Productivity. Review of Economics and Statistics. 2001/08/01 2001;83(3):434-445.
5. Rasel F. Combining Information Technology and Decentralized Workplace Organization: SMEs versus Larger Firms. *International Journal of the Economics of Business.* 2016:1-43.
6. Davenport TH, Short JE. The new industrial engineering: information technology and business process redesign. Sloan Management Review. 1990;31(4):11-27.

7. Allan M, Chisholm CU. Achieving engineering competencies in the global information society through the integration of on-campus and workplace environments. Industry and Higher Education. 2008;22(3):145-152.
8. Johri A. Situated engineering in the workplace. Engineering Studies. 2010/12/01 2010;2(3):151-152.
9. Todd RH, Sorensen CD, Magleby SP. Designing a senior capstone course to satisfy industrial customers. Journal of Engineering Education. 1993;82(2):92-100.
10. Todd RH, Magleby SP, Sorensen CD, Swan BR, Anthony DK. A survey of capstone engineering courses in North America. Journal of Engineering Education. 1995;84:165-174.
11. Miller RL, Olds BM. A Model Curriculum for a Capstone Course in Multidisciplinary Engineering Design. Journal of Engineering Education. 1994;83(4):311-316.
12. Huber GP. Organizational learning: The contributing processes and the literatures. Organization science. 1991;2(1):88-115.
13. Argote L, Epple D. Learning curves in manufacturing. Science. 1990;247(4945):920-924.
14. Argote L. Organizational learning: Creating, retaining, and transferring knowledge. 2nd ed. New York, NY: Springer; 2013.
15. https://products.office.com/en-US/sharepoint?legRedir=true&CorrelationId=6576d639-0145-43bf-8184-77ee91edc93c. Accessed last on January 28, 2016.
16. http://www.blackboard.com. Accessed last on January 28, 2016.
17. http://www.edx.com. Accessed last on January 28, 2016.
18. Madhavan K, Richey MC. Problems in Big Data Analytics in Learning. Journal of Engineering Education. 2016;105(1):6-14.
19. Hanneman RA, Riddle M. Introduction to social network methods. Riverside, CA: Published digitally at http://www.faculty.ucr.edu/~hanneman/nettext/; 2005.
20. Blumberg R, Atre S. The problem with unstructured data. DM REVIEW. 2003;13(42-49):62.
21. https://www.arangodb.com. Accessed last on Janueary 28, 2016.
22. http://www.tableau.com. Accessed last on February 1, 2016.
23. https://www.rstudio.com. Accessed last on February 1, 2016.
24. National Academy of Engineering. Grand challenges for engineering in the 21st century. 2008; http://www.engineeringchallenges.org/. Accessed March 9, 2011, 2010.
25. Jain AK. Data clustering: 50 years beyond K-means. Pattern recognition letters. 2010;31(8):651-666.
26. Murtagh F. A survey of recent advances in hierarchical clustering algorithms. The Computer Journal. 1983;26(4):354-359.
27. Lawton D, John B, Nancy V, Michael CR, Vivian TD, David EF. Learning science principles for effective online learning in the workplace. Proceedings of the IEEE Frontiers in Education Conference (FIE), 2010: F2D-1-F2D-6. doi: 10.1109/FIE.2010.5673309.