PROCEEDINGS
OF THE
THIRTEENTH ANNUAL ACQUISITION RESEARCH SYMPOSIUM

THURSDAY SESSIONS
VOLUME II

Big Data Analysis of Contractor Performance Information for Services Acquisition in DoD: A Proof of Concept

Uday Apte, Professor, NPS
Rene Rendon, Associate Professor, NPS
Mike Dixon, Assistant Professor of Operations Management, Ivey Business School

Published April 30, 2016

ACQUISITION RESEARCH PROGRAM
GRADUATE SCHOOL OF BUSINESS & PUBLIC POLICY
NAVAL POSTGRADUATE SCHOOL

Panel 20. Analyzing and Incentivizing Performance Outcomes

Thursday, May 5, 2016	

3:30 p.m. – 5:00 p.m.	**Chair: Mark Deskins,** Director, Acquisition Career Management, ASN(RD&A)
	
	 Big Data Analysis of Contractor Performance Information for Services Acquisition in DoD: A Proof of Concept
	
	      Uday Apte, Professor, NPS
	      Rene Rendon, Associate Professor, NPS
	      Mike Dixon, Assistant Professor of Operations Management, Ivey Business School
	
	 Knowledge, Experience, and Training in Incentive Contracting for the Department of Defense
	
	      Kevin Carman, Dean, DAU
	      Randall Gibson, Professor, DAU
	
	 Learning in the Shadows: A Retrospective View
	
	      Donna Kinnear-Seligman, Mission Assistance Program Analysis Manager, DAU

Big Data Analysis of Contractor Performance Information for Services Acquisition in DoD: A Proof of Concept

Uday Apte—is a Distinguished Professor of Operations Management at the Graduate School of Business and Public Policy (GSBPP), Naval Postgraduate School, Monterey, CA. He also serves as the Associate Dean of Research and Development at GSBPP. Before joining NPS, Dr. Apte taught at the Wharton School, University of Pennsylvania, Philadelphia, and at the Cox School of Business, Southern Methodist University, Dallas. He is experienced in teaching a range of operations management and management science courses in the executive and full-time MBA programs. Prior to his career in academia, Dr. Apte worked for over 10 years in managing operations and information systems in the financial services and utility industries. Since then he has consulted with several major U.S. corporations and international organizations.

Dr. Apte's research interests include service operations management, supply chain management, technology management, and globalization of information-intensive services. He has completed over 10 sponsored research projects for the U.S. Department of Defense and has published over 60 articles, five of which have won awards from professional societies. His research articles have been published in prestigious journals including *Management Science, Interfaces, Production and Operations Management (POM), Journal of Operations Management (JOM), Decision Sciences Journal (DSJ), IIE Transactions,* and *MIS Quarterly*. He has co-authored two books, *Manufacturing Automation* and *Managing in the Information Economy*. Dr. Apte has served as a vice president of colleges at Production and Operations Management Society (POMS), as a founder and president of the POMS College of Service Operations, and as guest editor of POM journal. Currently he serves as the senior editor of POM and as associate editor of DSJ.

Dr. Apte holds a PhD in Decision Sciences from the Wharton School, University of Pennsylvania. His earlier academic background includes an MBA from the Asian Institute of Management, Manila, Philippines, and Bachelor of Technology (Chemical Engineering) from the Indian Institute of Technology, Bombay, India. [umapte@nps.edu]

Rene G. Rendon—is an Associate Professor at the Graduate School of Business and Public Policy, NPS, where he teaches defense acquisition and contract management courses. He also serves as the Academic Associate for the MBA specialization in contract management. Prior to joining the NPS faculty, he served for over 20 years as an acquisition contracting officer in the United States Air Force. His career included assignments as a contracting officer for the Peacekeeper ICBM, Maverick Missile, and the F-22 Raptor. He was also a contracting squadron commander and the director of contracting for the Space Based Infrared Satellite program and the Evolved Expendable Launch Vehicle rocket program. Rene has published in the *Journal of Public Procurement*, the *Journal of Contract Management*, the *Journal of Purchasing and Supply Management*, and the *International Journal of Procurement Management*. [rgrendon@nps.edu]

Michael Dixon—is an Assistant Professor of Operations Management, Ivey Business School. He is interested in service operations management and specifically the effect of operational decisions and service design on customer experience. Dr. Dixon conducts management research from an analytical perspective and analyzes large archival data sets to model customer behavior, finding ways to represent the behavioral aspects of service design in quantitative and analytical ways. Specifically, he is interesting in understanding the role of experiential sequence on participant perception. In past research, he analyzed season subscription ticket sales for a large performing arts venue and showed that season subscription re-purchases are influenced by event sequences from previous seasons, leading to managerial direction for event planners to properly schedule performances across a season to maximize experiential effects. He received the Liskin Teaching Award from the Naval Postgraduate School in 2014 and 2015.

Dr. Dixon is a co-editor of the Production and Operations Management Society (POMS) Chronicle newsletter. He has publications in the *Journal of Operations Management* and articles under review at the *Production and Operations Management, Journal of Service Management*, and *Journal of Service Research*. He has served as a reviewer for *Management Science, Production and Operations*

Management, *Journal of Operations Management*, *Journal of Services Research*, *Service Science*, and the *Decision Sciences Journal of Innovative Education*. [mdixon@ivey.uwo.ca]

Abstract

This paper explores the use of Big Data analytic techniques to explore and analyze large datasets that are used to capture information about DoD services acquisitions. We describe the burgeoning field of Big Data analytics, how it is used in the private sector, and how it could potentially be used in acquisition research. We test the application of Big Data analytic techniques by applying them to a dataset of CPARS (Contractor Performance Assessment Reporting System) ratings of acquired services, and we create predictive models that explore the causes of failed services contracts using three analytic techniques: logistic regression, decision tree analysis, and neural networks. The report concludes with recommendations for using Big Data analytic techniques in acquisition.

Introduction

In April 2015, the Under Secretary of Defense for Acquisition, Technology, and Logistics (USD[AT&L]) issued his implementation guidance for Better Buying Power (BBP) 3.0 with the theme Achieving Dominant Capabilities through Technical Excellence and Innovation. The purpose of the BBP 3.0 acquisition initiative is to strengthen the Department of Defense's (DoD's) efforts in innovation and technical excellence while also continuing the DoD's efforts to improve efficiency and productivity (USD[AT&L], 2015). One of the major components of BBP 3.0 is its emphasis on improving the tradecraft in acquisition of services. The implementation guidance focuses on strengthening the contract management function for installation level services, improving requirements definition in the services acquisition process, and improving the effectiveness and productivity of contracted engineering and technical services.

It is not surprising that the USD(AT&L) has focused on improving services acquisition in the DoD. Services contracting specifically, and contract management generally, have been identified as a "high risk" by the Government Accountability Office (GAO). Since 1992, the GAO has found that the DoD lacks an adequate number of trained acquisition and contract oversight personnel, uses ill-suited contract arrangements, and lacks a strategic approach for acquiring services (GAO, 2015). Additionally, the GAO has reported that the DoD lacks adequate data needed to inform its decision-making on services acquisition and contract management. The GAO has also stated that the DoD lacks established metrics to assess its progress in improving services acquisition, and that the DoD should leverage its acquisition data by developing baselines to identify trends, thereby enabling it to develop measurable goals and gain more insight into whether its initiatives are improving services acquisition.

The purpose of this research is to explore how the DoD can leverage acquisition data, specifically contractor performance information, in identifying drivers of success in services acquisition. Through the use of exploratory descriptive and predictive statistical models, we describe and uncover the drivers of low and high contractor performance scores. In uncovering and describing these drivers, we develop recommendations for cost-effective management of services acquisition. Furthermore, we perform additional statistical analysis to determine if there is any relationship between contractor performance assessment factors (quality, schedule, cost, business relations, and management of key personnel), service type, contract type, level of competition, and contract dollar value. In researching the relationships among these variables, we perform predictive-modeling-based statistical methodology appropriate for Big Data including predictive regression modeling, decision-tree analysis, and neural-network analysis to determine which variables—

contractor performance data, contract characteristics, and management approach—can be considered the drivers of success for services acquisition.

This research report is organized into five sections. This introductory section is followed by the second section which reviews our past research in services acquisition with a focus on investigations into contractor performance information and drivers of success in services acquisition. The third section provides a primer on the use of Big Data analytics and selected Big Data analysis tools. The fourth section provides the results of our Big Data analysis on contractor performance information and its relationship to drivers of success in services acquisition. We complete the report in the fifth section with conclusions and recommendations for using Big Data analysis in investigating success drivers in services acquisition.

Past Research

We have addressed the need for research in the increasingly important area of services acquisition by undertaking six sponsored research projects over the past several years. The first two research projects (Apte et al., 2006; Apte & Rendon, 2007) were exploratory in nature, aimed at understanding the types of services being acquired, the associated rates of growth in services acquisition, and the major challenges and opportunities present in the service supply chain.

The next two research projects were survey-based empirical studies aimed at developing a high-level understanding of how services acquisition is currently being managed at a wide range of Army, Navy, and Air Force installations (Apte, Apte, & Rendon, 2008, 2009). The analysis of survey data indicated that the current state of services acquisition management suffers from several deficiencies, including deficit billet and manning levels (which are further aggravated by insufficient training and the inexperience of acquisition personnel) and the lack of strong project-team and life-cycle approaches. Our research (Apte, Apte, & Rendon, 2010) also analyzed and compared the results of the primary data collected in two previous empirical studies involving Army, Navy, and Air Force contracting organizations so as to develop a more thorough and comprehensive understanding of how services acquisition is being managed within individual military departments.

As a result of these research projects dealing with the service supply chain in the DoD, we have developed a comprehensive, high-level understanding of services acquisition in the DoD, have identified several specific deficiencies, and have proposed a number of concrete recommendations for performance improvement.

In our research, we analyzed 715 Army Mission Installation Contracting Command (MICC) service contracts found in the PPIRS database. These contracts were specifically for professional and administrative, maintenance and repair, utilities and housekeeping, and automated data processing and telecommunication services (Hart, Stover, & Wilhite, 2013). The results of our analysis of contract variables and contract success (Rendon, Apte, & Dixon, 2014) are summarized as follows.

- Utilities and housekeeping services had the highest failure rate of all the product service codes analyzed. The reasons for contract failure included business relations and management of key personnel.

- Contracts with a dollar value from $50 million to $1 billion had the highest failure rate of all the contract categories. This group's most common reason for failing was cost control.

- Contracts awarded competitively had the highest failure rate when compared to the other two forms of competition available. The reasons that most often resulted in a contract failure were in the areas of schedule and cost control.

- Contracts structured as a combination contract type had the highest failure rate when compared to the other five types of available contracts.

In this past research (Rendon et al., 2014), we further analyzed our contract data to determine whether any of the variables had a significant relationship with contract success by specifically looking at the contract failure rates. We used the chi-square test (Fisher's exact test) to test if the actual failure rates are significantly different than what would be expected if the total contract failure rate was applied to each variable. The results of the chi-square test identified that Contractual Amounts and Contract Type were our only statistically significant variables.

We also looked at the relationships between percentage of filled 1102 billets and failure rates, and between workload dollars per filled billet and failure rates, and made some interesting observations. We saw that as the percentage of 1102 filled billets increased, the contract failure rate decreased. This would seem intuitive, that as the workforce increases, the contract success rate would also increase, since there would be sufficient resources to manage the contracting process.

In our most recent research (Rendon, Apte, & Dixon, 2015), using the original data set of 715 Army service contracts (Hart et al., 2013), we analyzed the narrative section of the CPARS (Contractor Performance Assessment Reporting System) reports to determine alignment with the objective assessment ratings (Black, Henley, & Clute, 2014). Based on interviews, we also analyzed the value added, not only of the narrative section, but also of the usefulness of the CPARS as a contractor assessment tool. Our focus was to recommend improvements to the CPARS contractor performance information documentation process. The results of our analysis of CPARS narratives and interviews, reported earlier in Black, Henley, and Clute (2014), are summarized as follows.

- The contracting professionals are doing a better job at providing beneficial CPARS data in the narrative when the contract is unsuccessful versus when it is successful.

- The contracting professionals were slightly better at matching the narrative sentiment to the objective scores in unsuccessful contracts than in successful contracts.

- The results of the interviews found that the CPARS database is still often not reliable, robust, or comprehensive enough. The interviews also reflected that unsuccessful contracts tend to have more reliable, robust, and comprehensive past performance information available in their CPARS reports. The interviewees also stated that the information found in the PPIRS database sometimes contains information in the narrative that is either contradictory or does not quite match up with the objective ratings.

In our current research, we use exploratory descriptive and predictive statistical models to describe and uncover the drivers of low and high contractor performance ratings. Additionally, we perform statistical analysis to determine if there is any relationship between CPARS factors and contract variables, as reflected in Figure 2. In researching the relationships among these variables, we perform predictive-modeling-based statistical methodology appropriate for Big Data including predictive regression modeling, decision-tree analysis, and neural-network analysis to determine which variables—CPARS factors, contract variables, characteristics, and management approach—can be considered as the

drivers of success for services acquisition. The next section of this report provides a primer on the use of Big Data analytics and the various types of Big Data analysis tools.

Big Data Analysis

The term *Big Data* is fairly new in modern business nomenclature. It refers to the massive influx of data that has been and is currently being collected in the digital and Internet era. In some estimates, 90% of the data that is currently being stored on computers and servers around the world was collected in just the past two years (Baesens, 2014, p. 1). Other authors (Mayer-Schoenberger & Cukier, 2013, p. 28) cite that in the year 2000, only one quarter of the world's data was digitized; the remainder was on paper and other analog media. However, by 2013, 98% of all data was digital.

The flood of data comes primarily from the digitization of processes, interactions, and communications brought about by digital innovations such as internet-consumerism, mobile technology, and social networking (Mayer-Schoenberger & Cukier, 2013). In addition, data storage capacity is becoming ever cheaper, making it easier to keep data indefinitely. The term *datafication* refers to turning aspects of life that, in the past, have never been quantified into data that can be analyzed; for example, GPS coordinates are being recorded in mobile transactions or photos, photo images are being "datafied" to find face matches by Facebook, and words and sentences from Twitter status updates are being analyzed for content and sentiment using various text analysis techniques.

The term *Big Data* is used to discuss how to store, manage, and—perhaps most importantly—analyze these large stocks of data. Specifically, Big Data analytics refers to the ability to make distinct observations from large amounts of data that might not be able to be inferred from smaller amounts (Mayer-Schoenberger & Cukier, 2013). According to these authors, Big Data analytics differ from traditional statistics in three important ways. First, sample sizes are much bigger, approaching at times the size of an entire population. Traditionally, statisticians use small, unbiased samples to make inferences about larger populations, which has worked well for simple questions. Complicated sampling techniques have to be deployed for more complex, layered questions in order to make inference about specific sub-groups of a population. Second, Big Data analytics have to settle with unclean data. Finally, Big Data analytics leads to correlational explanations and not causational, that is, the results of Big Data analytics can only be interpreted as correlational relationships between variables.

The new term *data science* refers to the skillset needed to make sense of Big Data (see Schutt & O'Neil, 2013). A data scientist is made up of equal parts computer scientist, statistician, mathematician, and graphic designer, with capabilities to pull and combine datasets; manipulate, clean, and analyze data; and communicate aggregate results in a meaningful way. Data scientists are found across multiple sectors, including journalism, academia, information technology, banking, insurance, sports, and government.

Big Data is used by computer scientists that feed computers volumes of data with hopes that computers can make inferences on the probability of intuitive analytics that, in the past, have proven very difficult to teach to a computer. The success of the IBM Watson project provides evidence that Big Data analytics can outperform the world's most clever trivia masters. Big data analytic techniques are being used to generate algorithms for computer learning, search engines, and risk management.

The focus of this paper is to describe, as a proof of concept, how Big Data analytic techniques could be used to further the understanding of successes and failures of the DoD and other federal service contracts. Using the CPARS data previously described, we

consider the range of analytics that could be used to expand the research and practice of service acquisitions.

Typical Form of Big Data

Datasets used for Big Data analytics are usually formed by taking multiple measurements of multiple cases. Data is organized in rows and columns. Data in the same row are all from the same case or observation, and the columns have the same measurement or variable for all cases. Typically, a dataset's size is described by the number of cases and its number of variables. One of the variables is an identification number that is unique for that individual case. There may also be other identification variables that can be used to describe the case's membership to some other category; for example, the zip code, state, unit, etc. Identification variables can be used to extract data from other sources, adding to the number of variables available for analytic modeling.

Analytical modeling is a term that describes various methods that specifically quantify relationships between variables using past data as an indicator of how relationships form and how they might exist in the future. In predictive analytics, analysts create models that attempt to explain relationships between a specific target variable (sometimes called a dependent variable) and any number of input or independent variables. Analytic modeling has two important tasks: to predict outcomes of future cases, and to quantify relationships between inputs and target variables. These two tasks are not always congruent; at times a model might be very good at predicting future cases while at the same time present a challenge in interpreting relationships found in the data.

In most cases, target variables are either continuous across a large scale (e.g., dollars, time, or distance) or categorical with just two categories, that is, binomial (e.g., defaulting on a loan, failing an assessment, or repurchasing of a product). Binomial target variables take the form of either "yes" or "no." Less common, but still available, is predictive modeling with categorical target variables with more than just two categories.

Predictive modeling uses probability and statistics to estimate relationships between variables. In traditional statistics, a sample of cases is used to make these estimations and the model is used to infer something about a larger or future population. Using larger samples sizes found in Big Data allows the analyst to compare a model's ability to predict and describe relationships with existing data; analysts will randomly select a percentage of cases to be withheld during the model building phases. After a model is proposed, an analyst will "validate" the model using the withheld dataset to see how it would perform using existing data. Having a "validation" dataset adds to the ability to use the model outside the sample that is used to create it.

Predictive analytic models, estimated using Big Data, can provide a good indication of how target variables can be predicted using other measurements of a case. Predictive models are used widely in situations in which there is a complex set of variables, some of which might be correlated to a target variable for part of the time. Take, for example credit scoring in which lending companies will use a predictive model to assess the risk that a borrower might default on a loan (binomial target variable). Creating models using data from past lenders, a portion of which defaulted, credit issuers can make decisions about whom to offer credit. The model might show that people who are young and have little income are at high risk of default. However, the quantifiable relationships that make up the model are entirely correlational and cannot be said to cause default; that is, being young with low income does not cause default. We stress this important point that predictive models are correlational and should not be used to describe causes of target variables.

Decision Tree Analysis

Decision tree analysis is a predictive analytics technique that attempts to identify and isolate portions of a dataset that seem to act in similar ways in regard to a target variable. Target variables can be binary, nominal, or continuous. The purpose of a decision tree analysis is to propose a set of rules that can be used to estimate or predict a target variable.

To begin decision tree analysis, the methodology first identifies the independent variable that most discriminates the target variable, that is, the one in which a separation will lead to the most divergent prediction of the target variable. This is done by considering what the typical target variable will be if the data is divided at points within the range of values of all the independent variables. Most software that conducts decision tree analysis will algorithmically consider all division across all independent variables, giving each divergent scores using one of various methods. The independent variable with the highest divergent score is usually chosen to be the first "branch" in the decision tree. The division of the data results in "nodes" that are further divided by other variables in the same manner, resulting in a tree in which the "root" is on the top and the "branches" go down. The final "nodes" are called "leaves" and give a prediction of the target variable for data that fits within the path that leads to it. What results is a fan-shaped visual depiction of simple decision-based models that can be used to predict the target variable. In addition to providing a prediction model, decision tree models also provide a good interpretation of how different values of independent variables impact a target variable.

Typically, the more branches in a tree, the better a model can predict target variables in a training dataset; analysts typically have to set rules about when to stop branching within the training dataset. However, it is often the case that only a few branches are appropriate for validation data. To combat overfitting, an analyst can "trim" the branches of the tree back to only those that contribute to the prediction of the target variable of the validation data.

Logistic Regression

The next method we discuss is modeling a binomial decision variable using regression techniques. Linear regression is taught in most college-level statistics courses. In traditional regression, an analyst will estimate a model predicting a continuous target variable using any number of both continuous and discrete independent variables. In decision tree analysis, the "model" resulted in a visual tree diagram that can be used to interpret and predict outcomes of cases; in regression the result of the modeling is a mathematical equation that can consider values of new case in order to predict the target. Traditional regression analysis is considered "linear" because the resulting mathematical model is in the form of a linear equation representing a line, or a multi-dimensional surface, that has slope and intercept. The equation of traditional linear regression analysis takes the following form:

$$\hat{y} = b_0 + b_1 x_1 + b_2 x_2 + \cdots b_n x_n \tag{1}$$

In Equation 1, the x_n are the values of each of the independent variables and collectively the equation can be used to predict the value of a target variable, $\hat{y}$. The "slope" portion of the equations are called "coefficients" and can be used to formally and explicitly describe relationships between independent and target variables. In the previous equation, the $b_1, b_2, \dots b_n$ are the coefficients that are estimated for each of the independent variables. The coefficients are "estimates" in the same way that the average of a sample is an estimate of the average of an entire population. Through independent hypothesis testing, a p value for each coefficient is calculated that can be used by analysts to determine if a coefficient significantly influences estimation of the target variable (recall that a low p value means that a coefficient is significant).

The traditional linear regression assumes that the target variable is continuous (e.g., temperature, weight, dollars) across a scale. When a target variable is binary (e.g., defaulting on a loan, failing an assessment, or repurchasing of a product), analysts use an extension of traditional linear regression called logistic regression. In logistic regression, the target variable takes on the binary form of zeros and ones, that is, the analyst assigns one of the two options to take the value of 1 and the other to take the value of 0. In traditional regression, the estimated model can be used to predict the actual values of the continuous target variable; in logistic regression, the equation will instead predict the *probability* that the case will take the value of 1 (instead of 0). The equation for a logistic regression takes the following form:

$$Prob(y = 1 \mid x_1, x_2, \dots x_{n1}) = \frac{1}{1+e^{-(b_0+b_1x_1+b_2x_2+\cdots b_nx_n)}} \qquad (2)$$

Equation 2 reads that the probability that the target variable y is equal to 1 given a set of independent variables $(x_1, x_2, \dots x_{n1})$ is equal to the fraction that has 1 as the numerator and $1 + e^{-(b_0+b_1x_1+b_2x_2+\cdots b_nx_n)}$ as the denominator. The form of the fraction ensures that the probability will be between 0 and 1 and the exponential function allows the traditional linear equation $(b_0 + b_1x_1 + b_2x_2 + \cdots b_nx_n)$ to be represented linearly even if the target variable is binomial. Using the past data, software packages use an algorithm called "maximum likelihood" to find the value of the coefficients that best fit the past data to the equation form.

Typically the interpretation of the coefficients $b_1, b_2, \dots b_n)$ are converted into "odds" or more precisely into "log odds." Odds are the ratio of probabilities; for binomial variables, odds can be represented as follows:

$$Odds\ (y = 1) = \frac{Prob(y=1)}{Prob(y=0)} \qquad (3)$$

Since we are dealing with binomial variables, this can be rewritten as follows:

$$Odds\ (y = 1) = \frac{Prob(y=1)}{1 - Prob(y=1)} \qquad (4)$$

Reformulating the previous regression equation model in terms of odds, we get the following:

$$\ln\left(\frac{P(y=1)\mid(x_1,x_2,\dots x_{n1})}{P(y=0)\mid(x_1,x_2,\dots x_{n1})}\right) = b_0 + b_1x_1 + b_2x_2 + \cdots b_nx_n \qquad (5)$$

The right-hand side of the reformulated equation now mimics the linear regression equation and is now linear in term of log odds. This reformulation is called a "logit transformation." In order to interpret the coefficients from a logistic regression, an analyst would typically calculate the exponent of the coefficient e^{b_n} and interpret it in terms of the original probability equation. For example if the exponent variable is above 1, say 1.8, you would say that the probability that the target variable would take the value of 1 will increase by 80% for every unit increase in the independent variable. If the exponent variable is below 1, say .80, you would say that the probability that the target variable will decrease by 20% for every unit increase in the independent variable.

Just like in decision tree analysis, regression models can be "overfit" by including too many non-generalizable independent variables. In addition, analysts using regression methodologies need to be aware that when independent variables are highly correlated with one another, the interpretation of the model is called into question (this problem is called multicollinearity). Deciding which variables to use in a model is typically done in one of two ways: (1) independent variables are chosen based on preconceived or theoretical

understanding, or of their relationship with the target variable; or (2) independent variables are considered algorithmically to determine their individual contribution to an overall model. This algorithmic consideration of independent variables is typically known as "step-wise" regression and consists of calculating the "goodness-of-fit" for models with differing combination of possible independent variables. The model that can explain the most amount of the variation of the target variable with the least amount of independent variables is usually chosen because of its "parsimonious" appeal, that is, its ability to explain with little complication.

Neural Networks

The final type of data analytics technique that we evaluate in this research is neural networks. Neural networks gets its name from neural pathways and connection in brains; the way ideas, thoughts, and facts are connected together in a dense web of connections within the brain. These pathways often have nodes that act as connectors between disparate paths. In neural networking with Big Data, algorithms are deployed to uncover layers of connecting nodes between different independent variables in order to better predict the target variable.

Neural networks essentially involves creating a series of regressions to uncover hidden connecting nodes which are in turn used as input for additional regressions to find deeper connecting layers, eventually leading to a regression model of a prediction of a target variable. In short, it is a series of regression models uncovering latent connecting layers of data that can, in turn, be used to better predict target variables. Analysts can control the level of connecting layers and which independent variables to use in the initial phases. The end result is a prediction model that can be verified using an independent validation dataset. The logical structure of a neural networks model with a single hidden layer is shown in Figure 1.
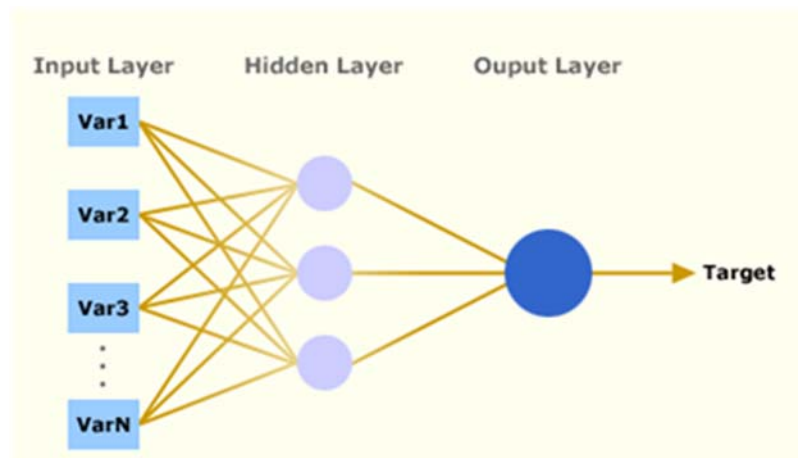

Figure 1. **Logical Structure of Neural Networks Model**

As we discussed in the previous section, regression techniques generally force analysts to create "linear" models, but using neural networks, analysts are able to model complex, nonlinear relationships using the intermediate layer nodes. The hidden layer nodes are able to handle the complexity of conditional (if/then) modeling that is not possible using traditional regression techniques.

Neural networks tend to work well with large datasets for which the analyst has very little preconceived theoretical model in mind. The results of a neural networks model are extremely difficult to interpret, and, as such, it is used primarily as a prediction modeling

technique as opposed to a descriptive or explanatory technique. Typically, the analyst is unable to describe the explicit connection between independent and dependent variables due to the complexities of the intermediate nodes.

Concluding Remarks

In addition to these methods, Big Data analysts are also concerned with topics such as missing data, data transformations, and model validations. Model validation will be addressed in subsequent discussions about training and validation datasets. Data transformations is a topic that is too broad for this paper, typically makes interpretation of results very challenging, and often leads to "overfitting" of the data. Missing data is often approached by "imputing" a value for data that is missing based on the mean or modes of the variable. In some cases, an analyst will infer a missing value based on a regression type formula with the missing value as the target variable. In our subsequent analysis, we imputed a small amount of missing data by replacing missing values with the mean value.

Big Data Analysis in Acquisition Research

A. Data Collection and Preparation

As mentioned earlier, the contract data used in our research was collected with the assistance of our graduate students (Hart et al., 2013). We searched the PPIRS database to identify Army Mission Installation Contracting Command (MICC) services (non-systems) contracts for the period 1996–2013. This search yielded 14,395 contracts in total. The data was then refined to include only those contracts associated with the following product/service codes:

- R: Professional, Administrative, and Management Support Services
- J: Maintenance, Repair, and Rebuilding of Equipment Services
- S: Utilities and Housekeeping Services
- D: Automatic Data Processing and Telecommunications Services

Based on the filtering for the previously mentioned service contracts, we identified 5,621 contracts. We then further filtered this database to include only contracts from the following Army MICC field directorate offices (FDOs) contracting organizations:

- MICC Region Fort Eustis
- MICC Region Fort Knox
- MICC Region Fort Hood
- MICC Region Fort Bragg
- MICC Region Fort Sam Houston

This data filtering resulted in 715 service contracts that were used in conducting our analysis, as seen in Table 1.

Table 1. Database Breakdown

(Hart et al., 2013)

	Total Contracts	
Total Army MICC Non-System Contracts	14395	
Less: Non R, J, S, D Service Contracts	8774	
Total R, J, S, D Service Contracts	5621	
Less: R, J, S, D Service Contracts at other MICC	4906	
R, J, S, D Service Contracts at MICC FDO Eustis, Knox, Hood, Bragg, Sam Houston	715	
Fort Eustis		238
Fort Knox		119
Fort Hood		114
Fort Bragg		55
Fort Sam Houston		189

For each contract, data was collected on specific contract variables (type of service, contract dollar value, level of competition, contract type) and specific contractor assessment ratings (quality of product/service, schedule, cost control, business relations, management of key personnel, and utilization of small business). Determining a contract to be successful or unsuccessful was made based on whether the contractor received a marginal or unsatisfactory rating in any of the CPARS assessment areas (quality of product/service, schedule, cost control, business relations, management of key personnel, or utilization of small business). The contractor receiving a marginal or unsatisfactory rating in any one of these assessment areas results in the determination of the contract as unsuccessful. It should be noted that the data collected from the PPIRS database was sanitized by removing identifiable data such as contract number, contractor name, DUNS number, and place of performance.

In addition to the contractor performance information accessed from the PPIRS-RC database, we also collected MICC region organization demographic data (annual workload in dollars, annual workload in actions, number of 1102 billets authorized, and percent of 1102 billets filled; Hart et al., 2013). This data was also analyzed to determine if these organizational demographics were related to contract success.

During our research we were able to receive access to PPIRS query tool that allows users to look up CPARS records individually. Unfortunately, we were not able to gain access to the CPARS databases with PPIRS directly; instead, we were required to pull records one at a time in order to conduct research. As previously described, our research team was able to pull 715 CPARS records (cases). While this is not a "Big Data" dataset, we believe that the actual CPARS dataset stored in PPIRS in its entirety is indeed Big Data. To our knowledge, there has been little to no research into this dataset. Therefore, in this paper we

propose several techniques that could be used to gain information from the Big Data that is being recorded and stored by the federal acquisition community.

Because our dataset is fairly small in Big Data terms, the results of our analysis should not be construed as being conclusive or indicative of general trends. However, if we are able to gain access to more or all of the CPARS records, the same analytics that we explore in the remainder of this paper can be used to gain a rich understanding of the dynamic and complex relationships between contracting attributes and CPARS scores. We intend to petition the gatekeepers of the CPARS records to make available the entire dataset so as to go forward with improved analytics.

In the following sections, we focus on three predictive modeling techniques: decision tree analysis, logistic regression, and neural networks. Each of these techniques has unique strengths to help researchers understand underlying relationships. All three are predictive modeling techniques that create models to predict a target variable. In our case, we use the CPARS data that we had collected for the previous studies; we use as a target variable a binomial indication of contract failure as previously described (a contract with either a marginal or unsatisfactory rating in any of the CPARS assessment areas.) As possible input variables we use the following variables:

- MICC
- Contract Start Month
- Contract Start Day
- Contract Start Year
- Contract End Month
- Contract End Day
- Contract End Year
- Fiscal Year of Contract
- Duration in days
- Contract Type: RJSD
- Awarded Dollar Value
- Current Dollar Value (at time of CPARS)
- Basis of Award
- Type of Contract (FFP, CPFF, CPAF, etc.)
- Annual Workload of Contracting Office (Dollars)
- Annual Workload of Contracting Office (actions)
- # of 1102 Billets Filled by Contracting Office
- % of 1102 Billets Filled by Contracting Office
- Workload ($) by Filled Billet
- Workload (actions) by Filled Billet

All analysis done in the following section was conducted using SAS Enterprise Miner, a leading software for Big Data analysis.

The first step in conducting any of the three types of analysis is to divide the original dataset into two datasets, the first being called a "training" dataset and the second called a "validation" dataset. The training dataset is used to create the analytical model, while the validation data is used to determine if the model is "overfit," that is, if the model is too

dependent on the training dataset to be applicable to other data. The validation data then "validates" the model that was created using the training dataset. Overfitting is a problem if the model is going to be used to predict target variables from observations outside what was used in the training dataset. In our case, we specified that 80% of the 715 cases be used for training the model and 20% be used to validate the model. The same cases were used to train and validate in all three techniques subsequently described.

Proof of Concept—Decision Tree Analysis

As discussed earlier, decision tree analysis is a predictive analytics technique that attempts to identify and isolate portions of a dataset that seem to act in similar ways in regard to a target variable. Figure 2 shows a decision tree we identified using SAS Enterprise Miner software for the binary target variable "unsuccessful contract." At the highest node, we see that 2.98% of the training dataset contracts were unsuccessful (1 = unsuccessful, 0 = successful) and 3.45% of the validation data. The first division is by the continuous variable called "Awarded Dollar Value"; those contracts that were less than $90,698,261 in awarded dollar value (ADV) had a much smaller failure rates (1.95% in training dataset and 3.05% in validation) compared to those that had higher awarded dollar value (12.07% and 7.14%).

The thickness of the line in the chart displays where the majority of the data lie; 512 cases in the training dataset had less than $90.6 million ADV while only 58 cases had more than $90.6 million ADV. Because there are so few cases with ADV greater than $90.6 million, there is little reason to further divide this section; however, if more data were available, the decision tree could be much more complex.

For those contracts with ADV less than $90.6 million, the next division is the "Workload (Actions) by Filled Billet." The contracting offices with less than 74.5 workload actions by filled billets had much lower failure rates (0.99% training, 3.7% validation) than that for offices with higher workload actions by filled billets (5.66% training, 0% validation). This would suggest that contracting offices that are understaffed or overworked tend to have larger number of contracts with low CPARS scores. However, take note that the validation dataset does not follow the same direction as the training dataset, suggesting that the model is overfit. Having a model that is overfit this early in a decision tree model is a symptom of having a small initial sample size.
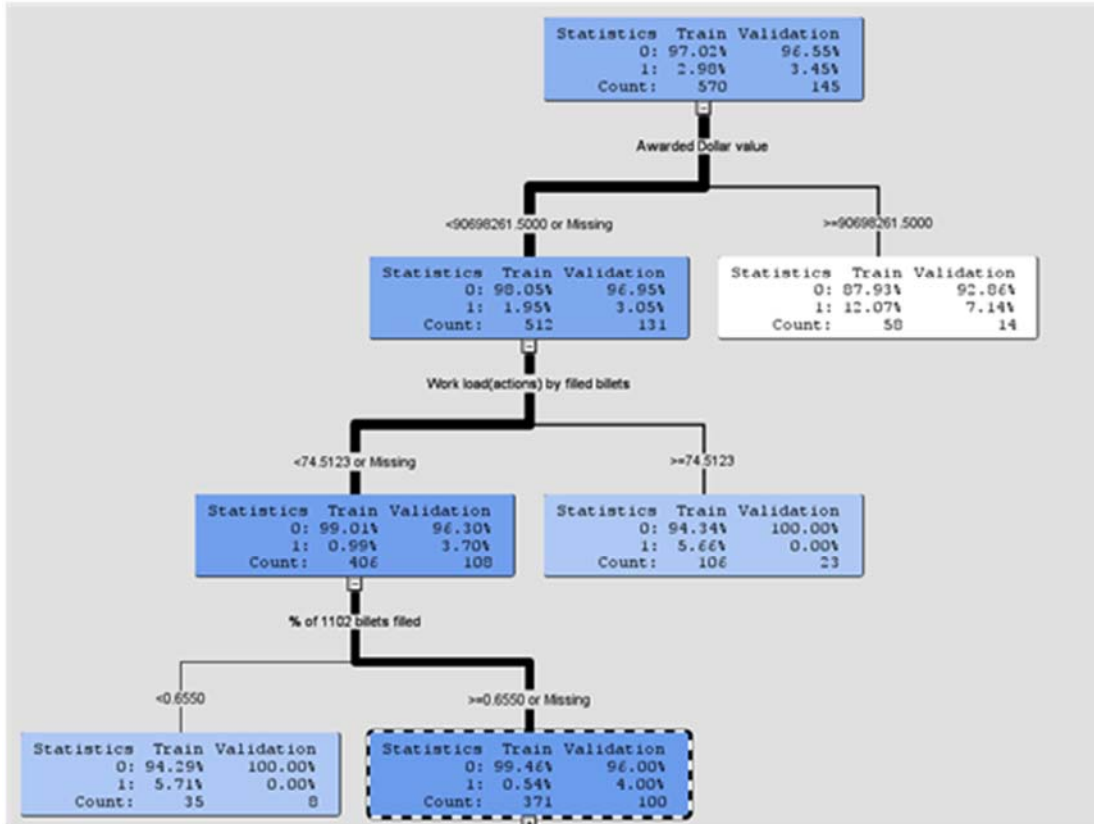
Figure 2. Decision Tree Analysis for "Unsuccessful Contract"

The final division happens with those contracts that are both less than $90.6 million ADV and from contracting offices with less than 74.4 workload actions per filled billet. The division shows that the offices that have less than 65.5% of their 1102 billets filled have a larger failure rate (5.71% and 0%) compared to those with a higher percentage of 1102 billets filled (0.54% and 4%). This suggests that contracting offices that are unable to fill their billets are likely to have higher rate of failed contracts.

Training Versus Validating

The decision tree presented in Figure 2 shows how the training dataset could best be divided into groups based on the independent variables. The resulting divisions make groups that are the most divergent in terms of the percentage of the binary target variable "unsuccessful contracts." Unfortunately, the "validation" dataset does not always follow the divergent nature of the training dataset, and, as a result, it appears that this analysis is overfit. If a model is overfit, it is less useful to generalize to other observations. However, overfit models can be useful in interpreting past data. In our case, the dataset is relatively small and therefore it is not necessarily very representative of any large set of contracts. Consequently, it is difficult to make any definitive or generalizable observations. However, the purpose of this research is to assess how Big Data analytics can be used to gain better understanding the success of contracts and that purpose has been well served with this proof of concept study.

Proof of Concept—Logistic Regression

As described in the logistic regression section, we performed the regression analysis using a step-wise regression methodology. In this method, a regression was estimated first with no independent variables; that is, with only an intercept. Next, a model was estimated with an intercept and only one variable that could explain the most variability in the target variable. Next, a model with an intercept and two top variables was estimated. This process was continued until all the independent variables had been included in the analysis. At the conclusion of the modeling, the software program displays which of the models explains most of the variability in the target variable with the least amount of independent variables. The results are shown in Table 2.

Table 2. Results of Stepwise Logistic Regression

Parameter	Estimate	p value	$e^{(Estimate)}$
Intercept	-12.213	<.0001	0
Work load actions by filled billet	0.0129	0.0117	1.013
Type of Contract – CPAF	8.8507	<.0001	6979
Type of Contract – CPAF & CPFF	-3.2748	0.9986	0.038
Type of Contract – CPFF	9.2498	<.0001	10402
Type of Contract – CPFF FFP	37.0026	0.9954	1.7×10^{16}
Type of Contract – CPIF	-3.3486	0.9978	0.035
Type of Contract – FFP	7.8061	.	2455
Type of Contract - Other	-3.7514	0.9970	0.0264

	Training	Validation
Average Squared Error	0.0266	0.0290
Misclassification Rate	0.0281	0.0276

The numbers in the "Estimate" column are the estimated coefficients for the regression equation previously described. A p value less than 0.05 is typically considered significant. The final column is the exponent of the estimate; these are easier to interpret since the original coefficient is in terms of log odds. This model reveals that two main characteristics of the contract tend to do a fairly good job of classifying failures (see the misclassification rate for training and validation datasets around 2.8%). Introducing additional variables to this model did not significantly improve the estimates.

The variable "workload action by filled billets" is the number of work actions that the entire office did divided by the number of filled billets that a contracting office had during the time period. The calculation provides an average number of actions worked for each billet filled. The logistic regression results show that an increase of one more worked action per filled billet would increase the odds of a failed contract by 1.013 or 1.3%. That means that increased workload of 10 actions per billet would be 13% more likely to have a failed contract. This variable was also a significant indicator of failure in the decision tree analysis.

The type of contract is also a significant indicator of CPARS failures in our dataset. The variable "Type of Contract" is a categorical variable with multiple different categories, as follows:

- CPFF   Cost Plus Fixed Fee
- CPAF   Cost Plus Award Fee
- CPIF   Cost Plus Incentive Fee
- FFP   Firm Fixed Price
- Other   Other types of contracts

Using categorical variables in regression requires analysts to construct "dummy variables" for each category that take binary values 0 or 1. A dummy variable is created for all categories except for one category which is referred to as the "base case." The coefficients for the regression models should be interpreted in terms of the base case. In our example, the base case is FFP contract. The interpretation of the coefficients for these variables is as follows: CPAF contracts are 6,979 times more likely to have CPARS failures than the FFP contracts in our dataset. CPFF contracts are 10,402 times more likely to have failed CPARS than the FFP contracts. All other categories of contracts are not significantly different from the FFP contracts. Interestingly, these findings were not uncovered in either the decision tree analysis or the previous research we did with this dataset.

Proof of Concept—Neural Networks

In our earlier introduction of the neural networks technique, we stated that this technique tends to work best using very large data sets. In addition, we stated that the modeling of neural networks is primarily only useful for prediction with no meaningful ability to describe or explain relationships between independent and target variables. Instead, neural networks modeling is described in terms of its ability to correctly predict cases in the validation dataset.

Given that our dataset was rather small (only 512 cases in the training dataset), the results of neural network modeling were not much better than those for the logistic regression modeling. We found that by using a simple neural network model with only one layer of hidden nodes, we could create a model that would mimic both the average squared error and the misclassification rates found on Table 2 reporting on the previously mentioned logistic regression model. Our conclusion is that because our dataset was limited in size, a more complex modeling technique such as neural networks did not improve the prediction capacity. Hence, it would be better for an analyst to stay with the logistic regression model which is easier to interpret. However, if a large dataset were available, the neural networks modeling could have been useful for risk prediction.

Conclusions and Recommendations

Conclusions

In the previous section, we applied three Big Data analysis techniques—decision tree, logistics regression, and neural networks—to the CPARS data as proof of concept. As discussed earlier, we found that the following four variables exhibit the largest impact on the success/failure rates of contracts:

- Type of Contract (FFP, CPFF, CPAF, etc.)
- Awarded Dollar Value
- Workload (Actions) by Filled Billets

- % of 1102 Billets Filled by Contracting Office

As noted earlier, the size of the CPARS dataset that was available and used in this research was rather small, and as a result the previously mentioned conclusions cannot be unequivocally considered as being definitive. However, based on the results of our prior research and on work experience of one of the researchers as a contracting officer, we have every reason to believe the previously listed variables play important roles in affecting the success/failure rates of contracts.

Regarding the applicability and use of three Big Data analysis techniques tested in this research, we found that the first two techniques are scalable in a sense that although they are ideally suited for analyzing large datasets, they are also useful for analyzing datasets of limited size. In contrast, the neural networks technique is not likely to be particularly useful unless the dataset being analyzed is large in size.

Recommendations for Big Data Analysis Techniques in Acquisition

The current DoD acquisition community uses a number of disparate databases that capture specific acquisition and contracting data. Some databases consist of structured data while others consist of unstructured data (Rendon & Snider, 2014). Structured data are typically comprised of program data and contract data that can be mined through data mining techniques. For example, FPDS-NG provides pre-award summary data of contracts awarded by federal executive agencies. This database provides contract specific data such as contracting agency, contractor, type of contractor, federal supply class or service code, contract type, level of competition, contract dollar value, and so on. Additionally, the DoD's Selected Acquisition Report (SAR) provides post-award information to Congress such as cost, schedule, and performance data for major acquisition programs. The SAR reports are generally submitted on an annual basis and reflect changes from the previous report such as cost variances, changes in procurement quantities and changes in earned value management (EVM) metrics. Other sources of acquisition data include the Federal Business Opportunities (FEDBIZOPPS) website that contains contract solicitations (e.g., requests for proposals), industry conferences notices, and contract award notifications. Another source of acquisition data, specifically contractor performance data, is the already discussed Past Performance Information Retrieval System (PPIRS) that contains the contractor performance report cards known as the Contractor Performance Assessment Reports (CPARS).

The previously mentioned databases provide both pre-award (inputs) and post-award (outputs) sources of acquisition data. The optimum use of Big Data analysis would be to apply Big Data analysis techniques to both input and output acquisition data to explore any relationships between acquisition inputs and outputs. We propose the following recommendations for these types of Big Data analysis techniques in defense acquisition, as reflected in Figure 3.
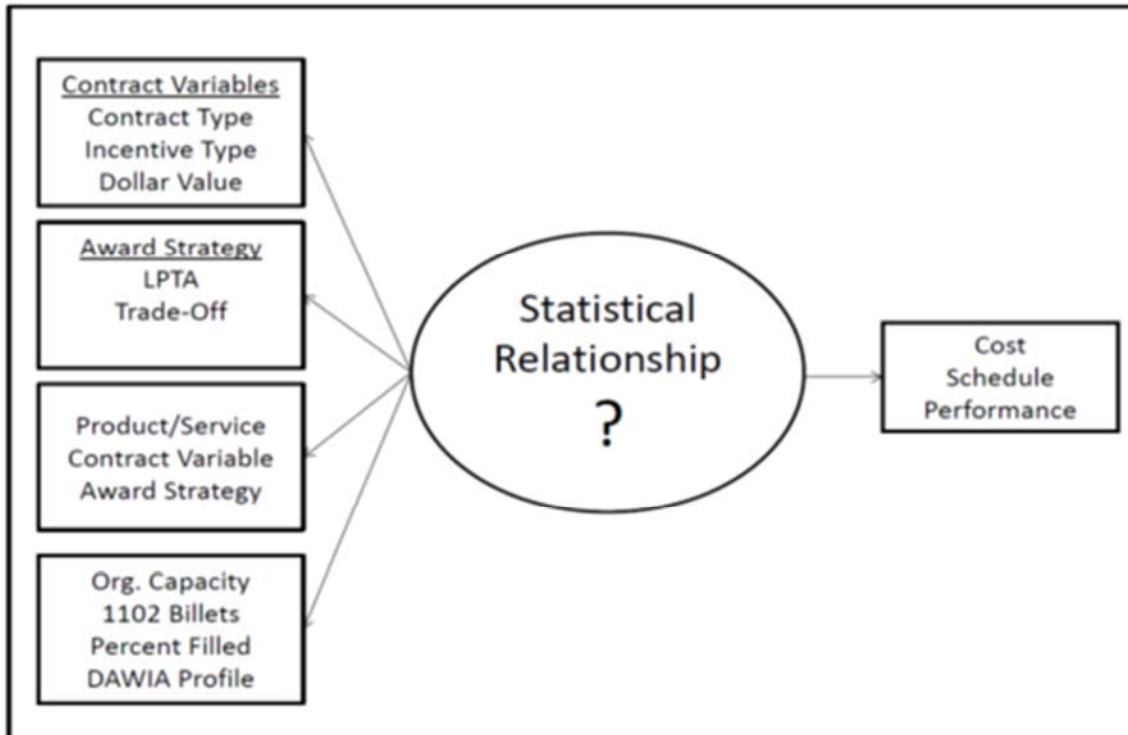
Figure 3. Proposed Recommendations

1. Analysis of specific contract variables and related contract cost, schedule, and performance outcomes. This Big Data analysis would look at the specific contract variables of contract type, incentive type, and contract dollar value and the resulting cost, schedule, and performance outputs of the contract. The purpose is to determine if contract type (fixed priced or cost reimbursement), incentive type (objective incentive such as FPI or CPI, subjective incentives such as award fee or award term), or dollar value is statistically related to the contract final cost, schedule, and performance results. This would require access and integration of the FPDS-NG, SAR, and PPIRS databases. The findings of this type of analysis would be beneficial in selecting contract type and incentive types on future contracts.

2. Analysis of specific contract award strategy variables and related contract cost, schedule, and performance outcomes. This Big Data analysis would look at the specific contract award strategy of price-based awards (such as lowest priced, technically acceptable) and tradeoff based awards (such as performance price tradeoff) and the resulting cost, schedule, and performance outputs of the contract. The purpose of this analysis is to determine if contract award strategy is statistically related to the contract final cost, schedule, and performance results. This would require access and integration of FEDBIZZOPPS database of solicitations, contract source selection files, SAR, and PPIRS databases. The findings of this type of analysis would be beneficial in selecting contract award strategies on future contracts.

3. Analysis of specific product/service codes, specific contract variables, contract award strategy variables and related contract cost, schedule, and performance outcomes. This Big Data analysis would look at the different

products and services procured by the DoD by product/service codes, as well as by contract type, contract award strategy and the resulting cost, schedule, and performance outputs of the contract. The purpose of this analysis is to determine if specific types of products or services are associated with specific contract variables and contract award strategy and if there is a statistical relationship with the contract final cost, schedule, and performance results. This would require access and integration of FEDBIZZOPPS database of solicitations, contract source selection files, SAR, and PPIRS databases. The findings of this type of analysis would be beneficial in selecting contract variables and contract award strategies on future procurement of specific products and services.

4. Analysis of organizational contracting capacity and related contract cost, schedule, and performance outcomes. Organizational contracting capacity includes metrics such as number of contracting (1102 and military equivalent) billets, percent of filled contracting billets, and number of DAWIA certified contracting personnel. This analysis would explore the relationship between the organization's capacity to contract (reflected in number and percent filled billets and DAWIA profile) and the organization's resulting cost, schedule, and performance outputs of its awarded contracts. The challenge in this Big Data analysis application is getting access to the organization's contracting capacity metrics. These metrics are not necessarily maintained by organizations, or may only be maintained at the higher headquarter levels. The benefit in conducting this Big Data analysis would be to see the relationship between contracting workforce (in terms of numbers and competence level) and contract performance.

References

Apte, A., Apte, U., & Rendon, R. (2008). M*anaging the services supply chain in the Department of Defense: An empirical study of current management practices* (NPS-AM-08-137). Retrieved from Naval Postgraduate School, Acquisition Research Program website: http://www.acquisitionresearch.net

Apte, A., Apte, U., & Rendon, R. (2009). *Managing the services supply chain in the Department of Defense: Empirical study of the current management practices in the Army* (NPS-AM-09-136). Retrieved from Naval Postgraduate School, Acquisition Research Program website: http://www.acquisitionresearch.net

Apte, U., Apte, A., & Rendon, R. (2010). *Services supply chain in the Department of Defense: A comparison and analysis of management practices in Army, Navy, and Air Force* (NPS-CM-10-161). Retrieved from Naval Postgraduate School, Acquisition Research Program website: http://www.acquisitionresearch.net

Apte, U., Ferrer, G., Lewis, I., & Rendon, R. (2006). *Managing the service supply chain in the U.S. Department of Defense: Opportunities and challenges* (NPS-AM-06-032). Retrieved from Naval Postgraduate School, Acquisition Research Program website: http://www.acquisitionresearch.net

Apte, U., & Rendon, R. (2007). *Managing the service supply chain in the US Department of Defense: Implications for the program management infrastructure* (NPS-PM-07-126). Retrieved from Naval Postgraduate School, Acquisition Research Program website: http://www.acquisitionresearch.net

Baesens, B. (2014). *Analytics in a big data world: The essential guide to data science and its applications.* Hoboken, NJ: John Wiley & Sons.

Black, S., Henley, J., & Clute, M. (2014, May). *Determining the value of Contractor Performance Assessment Reporting System (CPARS) narratives for the acquisition process* (NPS-CM-14-022). Monterey, CA: Naval Postgraduate School.

Cleland, D. I. (1986). Project stakeholder management. *Project Management Journal, 17*(4), 36–44.

Eisenhardt, K. M. (1989). Agency theory: An assessment and review. *Academy of Management Review, 14*(1), 57–74.

El-Gohary, N. M., Osman, H., & El-Diraby, T. E. (2006). Stakeholder management for public private partnerships. *International Journal of Project Management, 24,* 595–604.

Federal Acquisition Regulation (FAR), 48 C.F.R. ch. 1 (2015).

Fitzsimmons, J. A., & Fitzsimmons, M. J. (2006). *Service management: Operations, strategy, and information technology* (5th ed.). New York, NY: McGraw-Hill.

Freeman, R. E. (1984). *Strategic management: A stakeholder approach.* Boston, MA: Pitman.

GAO. (2002). Best practices: *Taking a strategic approach could improve DoD's acquisition of services* (GAO-02-230). Washington, DC: Author.

GAO. (2009). *Defense acquisitions: Actions needed to ensure value for service contracts* (GAO-09-643T). Washington, DC: Author.

GAO. (2015). *High risk: An update* (GAO-15-290). Washington, DC: Author.

Hagan, P., Spede, J., & Sutton, T. (2012, December). *Defining and measuring the success of services contracts in the United States Navy* (NPS-CM-12-201). Monterey, CA: Naval Postgraduate School.

Hart, J., Stover, A., & Wilhite, T. (2013, November). *Management levers that drive services contracting success* (NPS-CM-13-122). Monterey, CA: Naval Postgraduate School.

Lee, L., & Dobler, D. W. (1971). *Purchasing and materials management: Text and cases.* New York, NY: McGraw-Hill.

Luo, Y. (2002). Contract, cooperation, and performance in international joint ventures. *Strategic Management Journal, 23*(10), 903–19.

Mayer-Schönberger, V., & Cukier, K. (2013). *Big data: A revolution that will transform how we live, work, and think.* New York, NY: Houghton Mifflin Harcourt.

Moe, T. M. (1984). The new economics of organization. *American Journal of Political Science,* 739–77.

Nash, R. C., Schooner, S. L., O'Brien-Debakey, K. R., & Edwards, V. J. (2007). *The government contracts reference book: A comprehensive guide to the language of procurement.* Washington, DC: George Washington University.

Rendon, R. G. (2010, June). *Assessment of Army Contracting Command's contract management processes* (NPS-CM-10-154). Monterey, CA: Naval Postgraduate School.

Rendon, R. G., Apte, U., & Dixon, M. (2014, February). *Services supply chain in the Department of Defense: Drivers of success in services acquisition* (NPS-CM-14-001). Monterey, CA: Naval Postgraduate School.

Rendon, R. G., Apte, U., & Dixon, M. (2015, March). *Services acquisition in the Department of Defense: Analysis of operational and performance data to identify drivers of success* (NPS-CM-14-190). Monterey, CA: Naval Postgraduate School.

Rendon, R. G., & Snider, K. F. (Eds.). (2008). *Management of defense acquisition projects.* Reston, VA: American Institute of Aeronautics and Astronautics.

Rendon, R. G., & Snider, K. F. (2014). Retrieving what's already there: Archival data for research in defense acquisition. In J. Soeters, P. M. Shields, & S. Rietjens (Eds.), *Routledge handbook of research methods in military studies* (pp. 78–91). New York, NY: Routledge.

Seifert, A. J., & Ermoshkin, I. K. (2010). *Analysis of Government Accountability Office and Department of Defense Inspector General reports and commercial sources on service contracts* (MBA professional report). Monterey, CA: Naval Postgraduate School.

Schutt, R., & O'Neil, C. (2013). *Doing data science: Straight talk from the frontline.* Sebastopol, CA: O'Reilly Media.

Under Secretary of Defense for Acquisition, Technology, and Logistics (USD[AT&L]). (2010, June 28). Bet*ter buying power: Mandate for restoring affordability in defense spending* [Memorandum]. Washington, DC: Author.

Williamson, O. E. (1979). Transaction-cost economics: The governance of contractual relations. *The Journal of Law and Economics, 22*(3), 233–61.