STP 8-A-05 AI GROUNDWORK FINAL REPORT

Defense Logistics Agency R&D Weapons System Sustainment Program March 2019

Table of Contents

INTRODUCTION	1
TASK 1: CREATE A FOUNDATION OF DATA SCIENCE KNOWLEDGE	
MANAGEMENT	2
TASK 2: CREATE NOTIONAL ARTIFICIAL INTELLIGENCE MODELS	4
TASK 3: CREATE AN AI USE CASE PLAYBOOK	7
TASK 4: RESEARCH USE OF DLA AZURE CLOUD FOR DEVELOPMENT OF	FAI
BUSINESS CAPBAILITIES	9
TASK 5: DEVELOP AI POLICY RECOMMENDATIONS	11
TASK 6: INTEGRATE RESULTS OF TASKS 1-5 WITH DLA STAKEHOLDER	S 13

INTRODUCTION

DLA Leadership recognized the potential value of AI to the Agency's Mission and wants to invest in Al-powered systems to improve its operations in support of the Warfighter. The Weapons System Sustainment R&D Program (WSSP), led by Mr. Manny Vengua, launched the Artificial Intelligence (AI) Groundwork Project (STP 08-A-05) in FY19 to research and recommend strategies for enabling the adoption of AI across the Enterprise.

Prior to this project, AI and machine learning (ML) efforts at DLA have been siloed and therefore not aligned to a greater mission and purpose. When WSSP added multiple short-term AI prototype projects to its strategic vision, it decided to partner with Accenture to initiate a foundational project to unify the approach, policy, technical prerequisites, workforce upskilling, and broader DoD alignment.

In addition to the R&D Office, the Accenture team performed this work in coordination with key DLA stakeholders from the Analytics Center of Excellence (ACE) and the Office of the Chief Technology Officer. The project team developed a series of assessments, recommendations, and content to establish a strategic framework for DLA's journey to production-scale AI systems. However, this project is only the first step toward AI transformation. Execution and implementation of these recommendations across the Enterprise are required to empower its workforce, foster a modern AI culture, and drive better business outcomes through AI-powered solutions. The AI Groundwork Project has six key task areas:

- Task 1: Create a foundation for Data Science Knowledge Management
- Task 2: Create Notional Artificial Intelligence Models
- Task 3: Create an AI Use Case Playbook
- Task 4: Research of use of the DLA Azure Cloud for AI development
- Task 5: Develop AI Policy recommendations for AI Governance and Controls, Ethics, Workforce AI Education, AI Model Bias Avoidance, and Cybersecurity
- Integrate the results of Tasks 1-5 with DLA Data and Analytic Strategy, DLA Task 6: Enterprise Enablers Document and 2018 DoD AI Strategy documents

This Final Report contains summaries of the accomplishments, deliverables, and recommended next steps for each of the six task areas. It provides a high-level overview of the project as well as a catalog of the deliverables, which can be used to point Leadership to key documents and assets that Accenture has created for DLA.

Accenture recommends that DLA continue to invest in this holistic people/process/technology approach for developing production-grade AI capabilities. Furthermore, because DLA is taking early action, DLA should share both its successes and lessons learned with the Office of the Secretary of Defense (OSD) and the Pentagon's Joint Artificial Intelligence Center (JAIC). Sharing Al project outcomes will promote DoD collaboration and will establish DLA as an Al "pathfinder" for the Fourth Estate, whose agencies will face many similar challenges.

This document contains recommendations and strategies only. It is the responsibility of DLA Leadership to determine their relevancy, weigh their value, and ultimately act to implement.

TASK 1: CREATE A FOUNDATION FOR DATA SCIENCE KNOWLEDGE MANAGEMENT

Accenture developed initial AI knowledge content specific to the DLA Mission that the Agency can leverage to build a modern AI workforce.

Scope

Create a foundation for Data Science Knowledge Management:

- Define a knowledge management curriculum and approach
- Develop a process to update, maintain, and expand an AI Knowledge Library
- Create separate artificial intelligence overview materials for executives and mid-level managers

Accomplishments

Accenture created an AI training approach that includes a framework for workforce upskilling as well as a recommended process for maintaining and updating the training materials. To establish a foundation for the knowledge content, the Accenture team developed AI Awareness Training, AI Conversant Training for Executives, and AI Conversant Training for Supervisors. The specifics of each training are detailed below.

1. Al Awareness Training:

This training builds awareness of AI fundamentals. The training is designed for all DLA employees to gain an understanding of AI basics and how AI can be adopted at DLA. Specifically, it outlines what AI is, the different types of ML and their application, value creation through AI systems, and progress at DLA to date.

2. Conversant Training for Executives:

This training enables DLA Leadership to deepen their knowledge of AI and expand their role in AI development. Executives need to understand how to remove barriers to change and how to drive AI value. Specifically, this training discusses how to inspire, target, build, execute, control and trust AI systems with respect to four core enablers: data, people, infrastructure and process.

3. Conversant Training for Supervisors:

This training aims to teach Supervisors about Responsible AI and outlines the central role that Supervisors play in AI system adoption. Specifically, it discusses how to identify good use cases and drive action for AI solution development. It also outlines the key groups within the organization that the Supervisor must actively collaborate with in order to successfully adopt AI.

Deliverables Created

- Al Training Approach & Update Process.pptx
- Awareness Training.pptx
- Conversant Training Executive.pptx
- Conversant Training Supervisor.pptx

Recommended Next Steps

The deliverables created for Task 1 represent the first step toward upskilling DLA's workforce and promoting AI culture across the Agency. Accenture Recommends the following next steps:

- 1. Perform an AI Awareness assessment for the entire workforce.
- 2. Align AI training content to ongoing DLA knowledge management system modernization.
 - a. Record training sessions for wider distribution
 - b. Create training assessments to gauge effectiveness
- 3. Implement an AI training approach and content update processes.
- 4. Develop a plan to provide AI Awareness training to all DLA employees.
 - a. It is recommended that the entire DLA workforce completes this training in order to establish a foundational knowledge of AI across the Agency.
- 5. Provide AI Awareness training to all DLA Employees.
- 6. Target employees for the Executive or Supervisor AI conversant trainings.
 - a. These trainings are designed for a limited audience who will be directly impacted by AI solutions or need to drive investment to enable AI.
- 7. Provide AI Conversant training to Executives and Supervisors (where applicable)
- 8. Assess 3rd party training options for job-ready and expert training curriculum.

TASK 2: CREATE NOTIONAL ARTIFICIAL INTELLIGENCE MODELS

Accenture built six Al Notional Examples using 12 distinct ML model types to demonstrate the data science process that powers Al solution development.

Scope

Create Notional Artificial Intelligence Models that can be easily tested on any platform for evaluation. The notional model documentation shall include the entire AI modeling process:

- Notional problem statement
- Development and feature selection of labeled training, dev, and test data sets
- Methods used to include algorithm selection, bias avoidance, parameter tuning, variable importance, model validation, and the training process.
- Training results from a minimum of three (3) algorithms ranging from simple to complex
- Final model testing results
- Recommended methods for implementation

Accomplishments

Accenture created six notional AI examples – using 12 distinct machine learning models – which guide users through a standard data science workflow for AI/ML solution development. Each example defines the notional problem statement, model development process, methodology, training process, final testing results, and overall instructions for usage. The notional models are packaged together for easy storage and transfer. Each model includes a corresponding dataset, Python code in an interactive Jupyter Notebook (as well as an HTML read-only version), and a PowerPoint guide that explains the process and concepts. A total of 12 machine learning algorithms are utilized across the notional use cases, providing DLA with a rich set of examples for 4 common AI problem types – regression, classification, time series, and computer vision. Details of each example are provided below:

1. Example 1A – Regression (Simple Linear):

This example uses Excel to provide intuition on machine learning using a Simple Linear Regression model to predict the price of an Uber ride based on the expected trip mileage. The Example is meant to demonstrate that fundamental machine learning principles and concepts are both familiar and accessible to non-experts. Also included in the package is a Jupyter Notebook with the Python equivalent of the model built in Excel.

2. Example 1B – Regression (Multivariate)

Expanding on 1A, Example 1B introduces multivariate regression techniques to predict the price of a home. It introduces the end-to-end data science workflow that is common across all Notional Models, explores more sophisticated algorithms and techniques for evaluation, as well as the idea that AI/ML approaches may not always be sufficient to solve the initial business problem. Ultimately, the end of the development process results in a go/no-go decision based on model performance and the expected ROI of production implementation.

3. Example 2A – Classification (Binary):

Example 2A introduces the second fundamental type of supervised machine learning problems: classification. The Example explores binary classification using a model that is trained to predict whether a financial transaction is fraudulent or not. In addition, different model performance metrics are introduced, along with the concept that accuracy is not always the best metric (and in fact is not preferred for classification problems.) The Example closes with a discussion on tuning hyperparameters to optimize model performance for the specific business problem.

4. Example 2B – Classification (Multi-Class):

Example 2B generalizes binary classification to multiclass classification, where more than two target classes exist. The notional problem involves predicting wine quality. Three models are built, trained, and evaluated using the metrics introduced in 2A. Additionally, open-source tools are introduced for efficiently tuning model hyperparameters and quickly evaluating multiple models.

5. Example 3 – Time Series:

Example 3 introduces complex problem types and advanced machine learning model architectures. Three different models are built and evaluated to predict the daily maximum temperature based on atmospheric conditions. Advanced analysis techniques are required when the historical data used for model training has a time dependence, and Example 3 demonstrates the value of recurrent neural networks to tackle this type of problem.

6. Example 4 – Computer Vision:

Example 4 introduces another advanced topic: computer vision and deep neural network architectures. Beginning with simple feed forward networks and expanding into state-of-the-art convolution networks, the example classifies objects in pictures into 10 classes (airplanes, cars, trucks, ships, birds, cats, deer, dogs, frogs, horses). The code for this Notional Model requires GPU hardware for training, so end users should only run the interactive code in an AI environment (e.g., Cloud).

7. Platform Evaluation Model

The Platform Evaluation Model is a refactoring of Example 4, which leverages the publicly available CIFAR-10 <u>dataset</u>. The platform evaluation model is coded in a .py script (not an interactive Jupyter Notebook) for portability and can be used to evaluate the computational performance of any AI-ready environment.

Deliverables Created

All notional model content is provided in one zip file named DLA Al Notional Models.zip. The zip file includes the following:

- Notional Models Overview.pptx
- Example 1A Simple Linear Regression
- Example 1B Multivariate Regression
- Example 2A Binary Classification
- Example 2B Multiclass Classification
- Example 3 Time Series
- Example 4 Computer Vision
- Platform Evaluation

Each example is in a separate folder containing the following files:

- Example_X_Guide.pptx
- Example_X_Data.csv
- Example_X_Code.ipynb
- Example_X_Code.html

Example 1A includes an additional file, Example_1A_Sandbox.xlsx, to introduce machine learning concepts to a non-technical audience using Excel. The Platform Evaluation contains the .py script, required data, and a readme file. For all examples, we recommend beginning with the .html and .pptx files.

Recommended Next Steps

All six notional examples use supervised learning models. This class of machine learning is the most widely used for business solution development; however, it is not the only type of machine learning. Accenture recommends that DLA take the following next steps:

- 1. Provide Notional Model Examples to ACE and Major Subordinate Commands Business Process data analysts interested in learning about the development of ML models.
- 2. Create additional notional models to introduce other machine learning techniques:
 - Unsupervised learning
 - Semi-supervised learning
 - Dimensionality reduction

TASK 3: CREATE AN AI USE CASE PLAYBOOK

Accenture developed a recommended AI Use Case Selection methodology for DLA to begin targeted AI development projects.

Scope

Create an AI Use Case Playbook, to include the following:

- Use case selection criteria
- Technology prerequisites
- Process milestones that enable the transition from use case R&D to production implementation
- Score card identifying best business Returns on Investment from Artificial Intelligence deployment – identify DLA business processes that would have the greatest benefit from the application of artificial intelligence.
- Identify best strategy and methods to ensure effective scalability of future AI deployments

Accomplishments

Accenture created an AI Use Case Playbook divided into three chapters: AI Use Case Selection, Technology Prerequisites, and Process Milestones - R&D Pilot to Production. The Use Case Scorecard and DLA Business Process Areas to Target are subsections of the AI Use Case Selection chapter. Best practices for scaling AI solutions are included as a subsection to the Technology Prerequisites chapter. Additional details of each chapter are provided below:

1. Al Use Case Selection:

This chapter outlines a framework to generate, analyze, and prioritize AI use cases. The methodology includes five steps – Generation and Intake, Definition and Refinement, Feasibility Evaluation, Prioritization, and Selection. The Prioritization step includes a scorecard to rank the AI use cases and facilitate Selection. The chapter closes with Accenture's recommendations for DLA Process Areas that have the greatest potential benefit from AI.

2. Technology Prerequisites:

This chapter includes 5 subsections – Buy vs. Build decisions, Data Pipeline, Al Model Training Environment, Deployment Channels, and Best Practices to Scale AI. When developing AI systems, the first decision is whether to buy a commercial product or build a custom system. Accenture provides buy vs. build decision logic. Building an AI system requires data, model training infrastructure, and production deployment channels. Each of these three topics are covered with in-depth recommendations. Finally, as DLA expands its initial R&D efforts, best practices for scaling AI solutions are provided.

3. Process Milestones – R&D to Production:

This chapter describes two processes executed sequentially that cover key milestones for transitioning from R&D project to Production system. The first process describes the AI model development (R&D) including two Go/No-go reviews. The second process outlines the AI Application development process (Production) with two additional Go/No-go reviews. The nature of R&D and AI system development requires checkpoints to review progress before additional investment is authorized.

Deliverables Created

- DLA AI Use Case Playbook.docx
- DLA AI Use Case Scorecard.xlsx

Recommended Next Steps

The AI Use Case Playbook provides a recommended framework and activities to identify high value and high likelihood of success AI use cases. Accenture recommends DLA take the following next steps:

- 1. Implement the AI Use Case Selection Methodology.
- 2. Develop a pipeline of AI Use Cases for evaluation including Design Thinking sessions for individual process areas.
- 3. Continuously adjust the AI Prioritization scorecard as DLA generates new data from executed AI Use Cases.
- 4. Develop similar "Playbooks" for AI Model Development, AI Application Development, and AI Data & Technology to foster an end-to-end AI methodology.

TASK 4: RESEARCH USE OF DLA AZURE CLOUD FOR AI DEVELOPMENT

Accenture researched the use of DLA Azure capabilities to accelerate AI model development and documented recommendations.

Scope

Research of use of the DLA Azure cloud for development of AI business capabilities:

• Identify the infrastructural constraints of the DLA Azure cloud as applicable to the transfer of data for AI Use Cases

Accomplishments

Commercial clouds offer a variety of tools, infrastructure, and capabilities to accelerate the development of AI systems. Microsoft Azure created dedicated Government and DoD cloud regions designed to meet additional security requirements. Before utilizing any cloud services, DLA must understand Defense Information Systems Agency (DISA) FedRAMP+ and Impact Level security controls. An overview of the DLA Azure Assessment is detailed below:

1. DLA Azure Usage for AI:

Accenture reviewed the rationale for Azure usage in AI development and identified the advantages and disadvantages of cloud usage for AI workloads. A review of cloud service types is included for security discussion.

2. DLA Azure Infrastructural Constraints:

Accenture identified six infrastructural constraints of DLA Azure cloud: (1) the Security Paradigm, (2) Data Transfers, (3) Configuration Process, (4) Repeatability & Scalability, (5) Software Approval for Data Science, and (6) Azure DoD Region Service Availability.

3. Review of Azure's AI Capabilities:

Accenture provided an assessment of selected Azure AI capabilities including Virtual Machines, Azure AI development accelerators, and Azure Cognitive Services. All reviewed AI capabilities were aligned to availability in the DoD IL 4/5 region.

Deliverables Created

• DLA Azure Al Assessment.docx

Recommended Next Steps

1. Perform similar assessment of other Cloud Service Providers (CSP) usage for Al development.

9

- 2. Assess On-Prem technology infrastructure and approved software to develop Al capabilities including but not limited to SAS Grid, Hortonworks Apache Spark & Hadoop, Enterprise Data Warehouse, Director's Dashboard, Python, and R (if approved).
- 3. Develop robust and automated process to provision cloud services for AI.
- 4. Address process constraints impeding full Azure potential for AI.
- 5. Create AI Training Cost Model to estimate project cloud service costs.

TASK 5: DEVELOP AI POLICY RECOMMENDATIONS

Accenture provided a series of policy recommendations to create a robust, ethical, and secure AI culture at DLA.

Scope

Develop AI Policy recommendations that address the following:

- Al governance and controls
- Ethics
- Workforce AI education, inclusion and engagement
- Al model bias avoidance
- Cybersecurity

Accomplishments

Accenture compiled policy recommendations for DLA in each of the categories listed above. Further details for each category recommendation are provided below.

1. Al governance and controls:

This section introduces and recommends key components for successful Al implementation, including an Al Organizational Structure, an Al Organization Operating Model, a Responsible Al Framework and Responsible Al Controls. Further, it provides recommendations for complying with external regulatory actions.

2. Ethics:

This section identifies and addresses five common ethical issues with AI adoption and recommends seven best practices for mitigation.

3. Workforce AI education, inclusion and engagement:

This section demystifies workforce concerns for job displacement by recommending a collaborative future and an alliance between humans and machines.

4. Al model bias avoidance:

This section identifies where bias can exist within data, models, and operations. Additionally, it recommends how to avoid bias from the earliest stages of AI adoption and discusses available bias assessment tools to augment the capabilities of data scientists.

5. Cybersecurity:

This section introduces security procedures with regard to the demands of AI and DLA's cloud security controls. Specifically, Accenture recommends the use of cloud technology for AI development and the adoption of a Development Security Operations (DevSecOps) framework to drive AI innovation while maintaining robust security.

11

Deliverables Created

• Al Policy Recommendations.docx

Recommended Next Steps

Accenture recommends that DLA leverage these AI Policy Recommendations to transform itself into an AI-powered organization.

- 1. Create an AI organization within the Analytics Center of Excellence.
- 2. Institute AI Governance and Responsible AI framework that aligns practices to current AI regulatory proposals.
- 3. Collaborate with workers to co-create a human-machine alliance, including upskilling, new job roles, and adopting new channels to reach Al-powered users.
- 4. Develop an AI Fairness Tool to assess sources of Bias within AI systems and make model adjustments for bias correction.
- 5. Implement DevSecOps in order to accelerate AI model and AI application development, AI system maintenance, and scale AI solutions.
 - Note: DevSecOps has many benefits for overall IT application development and maintenance.

TASK 6: INTEGRATE RESULTS OF TASKS 1-5 WITH DLA STAKEHOLDERS

Accenture integrated all project deliverables with existing internal and external AI strategy to accelerate DLA's establishment as the AI Pathfinder for the Fourth Estate.

Scope

Work with DLA Stakeholders to integrate the results of Tasks 1-5 with:

- DLA Data and Analytic Strategy
- DLA Enterprise Enablers Document; Objective 3.5
- 2018 DoD AI Strategy

Accomplishments

Accenture reviewed the three documents listed above and created an integration approach for each document. The integration details are summarized in a single document for the R&D Office, which is designed such that recommendations for each document stand alone. Details of the potential integration of deliverables from this project with the three strategic documents are provided below:

1. DLA Data and Analytics Strategy:

The DLA Data and Analytic Strategy has the most areas of overlap and therefore the most integration points. Accenture met with Steven Borgesi (J67A / document owner) to review recommendations and receive feedback. Steve agreed with areas for integration, pending receipt of this project's deliverables. The scope of this project overlapped with 5 of 6 key lines of effort within the DLA Data and Analytics Strategy. Al is a subset of analytics and powered by data, thus all deliverables augment the existing lines of effort with details specific to Al solution development.

2. DLA Enterprise Enablers Document; Objective 3.5:

The DLA Enterprise Enablers Document outlines objectives and strategies for driving innovation, data management, and technology change at DLA. The R&D AI Groundwork project deliverables are a first step toward accomplishing three objectives, including the innovation methodology, Objective 2.3 (Conduct data science activities), and Objective 3.5 (Use artificial intelligence to increase automation, efficiency and quality). Accenture is providing materials and recommendations that will accelerate the execution of these objectives.

3. 2018 DoD Artificial Intelligence Strategy:

The 2018 Department of Defense Artificial Intelligence Strategy provides a high-level vision of AI applications across the DoD and defines core enablers. DLA's on-going investments in AI, including this R&D project, are aligned to the broader DoD goals. The R&D AI Groundwork project deliverables provide DLA with initial actions and recommendations for expanding the use of AI in accordance with the DoD strategy. As such, DLA should utilize DoD investments as well as continue its current AI journey.

Deliverables Created

• R&D AI Integration.docx

Recommended Next Steps

The R&D AI Groundwork project creates multiple assets that DLA can utilize as part of its AI transformation journey. Accenture recommends the following next steps:

- 1. Build upon the initial foundation for AI implementation to continue DLA's AI transformation
- 2. Collaborate with the Joint Artificial Intelligence Center and greater OSD to establish DLA as a "pathfinder" for AI logistics support, including the broader Fourth Estate.
- 3. Align AI strategy and approach with Data and Analytics Strategy and DLA Data and Analytics Governance Plan.
- 4. Utilize deliverables that create the foundation to accomplish Enterprise Enablers Objectives and support target user groups to conduct data science activities and develop AI systems.