PacificCPE

COURSE
82402A

Managing Data Analytics and Artificial Intelligence

Data is only valuable when it's well-managed and properly governed. This course covers the frameworks and practices needed to oversee data analytics programs, then transitions into one of today's most talked-about technologies: artificial intelligence. You'll build a working understanding of what AI is, how it functions, and why it matters for accounting and tax professionals. If you want to move beyond the buzzwords and develop practical knowledge you can apply in your work, this course is a strong starting point.

CREDIT HOURS 2 | FIELD OF STUDY BUSINESS | EXAM QUESTIONS 9

Start Course ↓ | Final Exam Copy ↓

LEARNING ASSIGNMENTS AND OBJECTIVES

Assignment

- Study the full course materials included here.
- Complete the review questions after each chapter.
- Answer 9 final exam questions.
- Final exam is scored immediately, 70% or better passes.
- Certificate of completion is available in your account.
- Failed exams can be retaken free of charge.

Course Objectives

- Identify the characteristics of quality data.
- Recognize the characteristics of various types of AI algorithms.

Copyright Notice

© 2026 Pacific CPE, LP. All rights reserved.

This course and test have been adapted from supplemental materials and uses the materials entitled Emerging Technologies for Business Professionals © 2024 by Nishani Vincent and Amy Igou. Displayed by permission of the publisher, John Wiley & Sons, Inc., Hoboken, New Jersey. All rights reserved.

The materials in this course — including the text, illustrations, exhibits, examples, exam questions, and review questions — are the copyrighted property of Pacific CPE, LP. No portion of this course may be reproduced, stored in a retrieval system, distributed, or transmitted in any form or by any means (electronic, mechanical, photocopying, recording, or otherwise) without the prior written permission of Pacific CPE, LP, except for the personal, non-commercial use of the registered learner.

Pacific CPE, LP grants the registered learner a single-seat, non-transferable license to access this course for the purpose of completing the associated continuing professional education credit. Account credentials may not be shared.

This course is provided "as is" for educational purposes. While Pacific CPE, LP makes a reasonable effort to ensure the content is current and accurate, the materials should not be relied upon as a substitute for professional advice. Pacific CPE, LP disclaims all liability for any error, omission, or loss arising from the use of, or reliance on, the course content.

Publication Date: August 2026

TABLE OF CONTENTS

CHAPTER 01

Data Analytics Governance And Management 4

1.01 Data Analytics Projects 5

1.02 Implementing Analytics 5

1.03 Data Analytic Process 8

1.04 Data Quality 11

1.05 Data Governance 14

1.06 Ethics 17

Test Your Knowledge 22

CHAPTER 02

Introduction To Artificial Intelligence 25

2.01 Artificial Intelligence 26

2.02 How To Build An Ai Model 27

2.03 Types Of Artificial Intelligence 28

2.04 Ai Platforms And Software Applications 34

2.05 Use Cases 35

Test Your Knowledge 39

Final Exam Copy 59

Glossary 42

Index .....

Chapter 1: DATA ANALYTICS GOVERNANCE AND MANAGEMENT

Chapter Objective

After completing this chapter, you should be able to:

- Identify the characteristics of quality data.

An excerpt of Mathew Chacko's, Coca-Cola's director of enterprise architecture, interview with Sam Ransbotham, an associate professor of information systems at the Carroll School of Management at Boston College and the MIT Sloan Management Review guest editor for the Data and Analytics Big Idea Initiative is given below.

Coke is the world's largest beverage company, with more than 500 brands and 3,500 products sold worldwide. In 2013, the company had $46.9 billion in net operating revenues, and a net income of $8.6 billion. It has about 250 bottling partners with 900 bottling plants, and employs over 700,000 system associates worldwide.

Can you talk a little bit about the challenges of marrying all these external datasets from your bottlers? That seems like a particular challenge at Coke, where there's this huge infrastructure of "data middlemen," for lack of a better term.

Chacko: Well, point-of-sale data, scan data, is actually very big. Our commercial department really wants to have good information about that from one set of customers so that they can take that information and go to other customers and say, "Hey, we have this other customer that did these kinds of campaigns with this kind of strategy, and they were able to see this lift."

Our team wants to go to our customers with real numbers. It's important for us when we go to our customers to be able to give them fact-based information. We also want to take in things like event data and social media data and provide these value services back to our customers. That's a big thing for us.

In some areas we have to have flexibility, and in other areas we can standardize. Currently our bottlers run their own system, and so they send us data in all sorts of different formats. We have to be flexible in being able to inject data. But when we transform that data, we need to transform into those standardized taxonomies or hierarchies.

We also have the reverse problem because we need to transform information back into the bottler's view – we have to give them back their information in formats they can read. We aspire to provide data as a service, both to our customers and to our bottlers.[1]

DATA ANALYTICS PROJECTS

Implementing an enterprise-wide data analytic project with many data sources can be challenging. However, the following best practice guidelines may help data analytics projects go as smoothly as possible.

- **Have a clear objective**. Determine what the overall goal for undertaking an analytics project is. Align the project goals with the business strategy.
- **Know your stakeholder needs**. The objective of analytics is to provide useful information for decision-making. Therefore, it is imperative that you understand the competing interests of different stakeholders. Once stakeholder needs have been identified, data analytic projects can focus on creating value while balancing these competing needs.
- **Know your technology**. Adopting the right technology tools is very important to get various datasets that are in silos connected.
- **Train your employees**. Hiring competent staff in data science is important. However, managers should also consider training their team on the functionality of the selected technology tools and business operations. Encourage your data scientists to work with messy data and trawl through the data until they find insights.
- **Ensure that action is possible**. Is it viable to take action on the findings of the analytics? According to a study by Gartner, only 20% of insights from data analytics will result in an action taken.[2] The remainder of the analyses will go unused. Without any possible action, the organization will not gain value.
- **Be flexible**. Generally, data analytics projects work better with an agile approach than a more rigorous development approach. Most of the time, the process will be more iterative and lead to the organization discovering additional use cases.
- **Take time to improve data quality**. Data quality issues can cause various problems, from inaccurate results to delayed projects. Taking time up-front to review and cleanse data will save significant time in the future and lead to better quality decisions.
- **Plan for the future**. Unlike other types of technical development, a business intelligence environment will continually change due to more data, changes in the business environment, and new information requirements for decision-making. It is better to have a plan in advance to manage changes.

IMPLEMENTING ANALYTICS

Even though analytics projects may be more iterative than other technical development projects, every analytic project may follow the software development lifecycle. A larger analytics project can encompass developing tools, such as a data warehouse, affecting data analytics projects throughout

the organization. The goal is to ensure data consistency while providing flexibility to explore the data with minimal constraints.

Planning

Even though you may be beginning an enterprise analytics project, other departments may have done some analytics work already. If that is the case, it would be helpful to start the new project by identifying any work that has been done previously in the various departments. Some of these may be scripts developed using end-user software tools. But getting a better understanding of what already exists will help centralize knowledge within the organization and streamline analytic projects in the future.

Next, determine the questions to answer with analytics. Prioritize which analyses to tackle first. Ask if the organization can take action due to this analysis or if it is informational only. The actionable analyses should be of higher priority as they will provide the most significant value to the organization. Therefore, create a process to communicate the results to the appropriate stakeholder and evaluate and monitor the outcomes of the analysis. The type of analyses conducted will help to decide the breadth of team members needed for the project. Be sure to include employees from each functional area who understand their department's data.

Develop an overall strategy for the organization's analytical needs for the next few years. Even though you may only focus on a few projects in the immediate future, many decisions that are made with the first vital projects will affect future projects. Consider whether data analytics will be centralized or decentralized, as this will determine the required controls and governance processes. The types of applications, oversight responsibilities, and reporting process will be determined by whether the organization uses a centralized, decentralized, or hybrid approach. Determine if the organization will select standardized tools to be used by all functions or if there will be some decentralized choices. If analytics is decentralized, establish a policy of which types of analysis end-users can perform and what will require oversight by analytics experts. Many organizations can use a combination of centralized and decentralized tools. For example, the organization can use a standard data lake and warehouse administration application for most of its central data storage. However, individual departments can choose an end-user analytics tool, such as Powerbase, Tableau, or others.

Analysis

Based on the questions generated in the planning phase, this phase determines what data is needed to perform the analyses. Evaluate whether the data exist internally or externally. Some data will likely exist internally, but it will probably reside in different systems and formats. Therefore, take the time to understand various internal data sources, where the data resides, the types of data formats, and who the data owners are.

Further, consider what external data is needed or would help to extend the analyses. Determine whether the data is accessible or if alternatives are available. Based on the list of data sources, determine what types of data storage are needed to allow the appropriate individuals to access the data. Evaluate if the amount of data requires a more extensive data solution, such as a data warehouse or a data lake.

If using a data warehouse, determine how much data cleansing will be required to make the data useable across the organization by various users. Develop a process to ensure that the data loaded into the data warehouse is valid and in the proper format.

Design

A significant part of the design will be modeling the data, especially if using a larger storage capacity, such as a data warehouse. The data model will provide a roadmap of all the data components for the warehouse, such as the databases, tables, and indexes. To complete this part of the project, you will need experts on data warehouse management with experience in data modeling and business experts with an understanding of the conceptual relationships between the data.

At this point, the analysts should design and implement the controls to manage the data. Design and implement formal processes to address questions such as the following:

- What data will be loaded into the data warehouse? How often?
- How will a new data source be approved?
- What types of end-user analysis tools will be available? Who will have access to those tools?
- What data can the end-users have access to?

In the future, the organization will most likely obtain new data sources, require additional analyses, or change business operations. Therefore, consider whether the system is scalable when designing your data warehouses or lakes. Looking ahead will reduce the probability that the system will need a major overhaul in a few years.

Implementation

In this step, you will implement the plan by loading the data and preparing the way for analysis work. Using the list of data requirements from the design and analysis phases, obtain the data to load into the data warehouse.

Next, clean and format the data to load into the data warehouse. Allow more time than expected for data preparation, as it will reduce issues in the future. Budget additional time for each new data source. Several items to consider when preparing the data are:

1. **Dates**. Are all dates in a consistent format and a form usable for calculations?

2. **Consistency**. Are the data in a specific field consistent throughout? For example, states can be consistently spelled out or abbreviated. Having a combination will affect data analysis and give inconsistent results or eliminate data if the analysis is based on the abbreviation, for example, rather than the full name of the state.
3. **Units of measure**. Do all data in a specific field have a consistent unit of measure? For example, are length measurements all in meters?
4. **Blank fields**. Is there missing data? Determine whether the rest of the record should remain or be removed.
5. **Irrelevant fields or records.** Data that will not be used can be removed from the dataset.
6. **Duplicate data**. After merging data from various sources, are there duplicate records to remove?

These are just a few examples of items to clean before using or loading the data into a data warehouse. After the initial data cleansing, perform some fundamental descriptive analysis on the data, such as central tendency statistics, a histogram, or an overall data profile, to verify the accuracy of the data. These simple analyses will help determine if your data is adequately cleaned and whether there are outliers.

Besides loading data, create user-specific dashboards, queries, or other reports during this phase to help end-users with specific analyses based on the planning stage.

Maintenance

Use the long-term strategy developed in the earlier phases to continue to improve and maintain the data storage, such as the data warehouse/data lake, data analytics applications, and the communication, reporting, and oversight processes. Appoint a team consisting of various functional disciplines to periodically evaluate changing business needs, the adequacy of data analytics to meet current requirements, and the effectiveness of the communication processes and report the findings to the steering committee.

DATA ANALYTIC PROCESS

When beginning an individual analytics project, it is helpful to break it into several tasks. The public accounting firm, EY, developed a four-step methodology to approach an analytics problem and create an analytics mindset. Even though there are other analytics processes available, they all follow similar steps, as listed below.[3]

The four steps a user can take when starting data analysis are:

1. Ask the right questions.
2. Extract, transform, and load the data.
3. Apply the appropriate data analytics techniques.

4. Interpret and share the results.

The process provides a roadmap for users to perform accurate analysis and ensure that the information generated is precise and reliable. These steps are further explained next.

Ask the Right Questions

Data analytics starts with determining what questions you are trying to answer. Users who attempt to skip this step often find themselves wandering in an attempt to find an answer. This step requires more thought than just jotting down some questions.

First, you need an understanding of the stakeholders/audience. Who will be the recipients of the data? What decisions are they trying to make? Understanding how the information will be used will help better define the scope of the questions required for the analysis.

Next, the analyst needs to understand the business environment well. Some analyses rely on industry knowledge to put the data into context. For example, you should understand the typical key performance indicators in the industry or any other norms to help compare the data. Besides understanding the industry, users must understand the business processes that generated the data. A deeper understanding of the context will generate better questions before the analysis.

Finally, determine how the analysis will be used.

- Can the organization take action or make a specific decision with this data?
- Will it be useful and provide value?
- Is the analysis required for external reporting purposes?

If the answer to these questions is no, reevaluate the purpose of the analysis.

Extract, Transform, and Load Data

The next step is to obtain and prepare the data for analysis. During this step, select which data analytic software applications you will use for the analysis. You may need one application to cleanse the data, another to perform the analysis, and another to visualize the data. The order in which data is extracted, transformed, and loaded will depend on the data analytics application(s) used.

Based on the questions generated in the previous step, you need to determine what data types are needed. The following is a short list of things to ask about each dataset.

- Where is the data stored?
- What date range of the data is needed?
- Is there a subset of data? For example, are all vendors needed for the analysis?
- Can I retrieve the data, or do I need to contact IT?

The answers to these questions may differ for each dataset. If this is an analysis that needs to be done regularly, establish a standard process or, if possible, automate the data collection, possibly with robotic process automation (RPA).

If the organization uses a data warehouse or data lake, then the data may already be transformed and loaded into a central repository. If this is the case, you will still need to determine which data tables must be accessed and understand the relationships between the data fields.

Once you determine the appropriate datasets, be sure you understand the data. Review the data dictionary to understand what each field in the dataset represents. Are any of the fields coded or linked to a master table that contains detailed information? For example, if the dataset has a column with a customer number, you may need a list of customers to provide the customer's name for the corresponding number.

Before starting your analysis, make sure the data has been cleaned. First, eliminate any data fields that are not required for the analysis. Next, review the data for data quality issues, such as inconsistencies in the data. The clean dataset reduces processing time later and allows you to focus only on the needed items. Once the data has been cleaned, it can be loaded into the data analytics tool.

Apply the Appropriate Analytics Technique

Applying the appropriate analytics technique can be an iterative process. Based on the questions in the first step and the data available, select the appropriate method for analysis. A high-level review of the basic techniques is in Chapter 3 of course 82401B.

When analyzing new data, analysts often start with the most basic form of analytics – descriptive. Descriptive analytics provides an overview of the data to confirm further plans for analysis. This analysis also determines if any additional data quality issues were missed in the previous step or if any outliers need to be examined.

Then additional analyses would progress through diagnostic and predictive methods, stopping when the question has been answered. Sometimes, the initial analysis choice does not provide the insights you wanted. Evaluate and determine if there is a more appropriate methodology, whether more data is needed, or whether the analysis is not possible now.

After the initial analyses, take a minute to review the results to ensure that they make sense within the context of the organization. Reviewing the findings before moving forward with an action plan ensures that the analysis is complete and accurate. Using your knowledge and a bit of skepticism, examine the results to determine whether they appear reasonable and explainable. With complex analysis, many things can alter the results, so performing a final review can be a good backup.

Interpret and Share the Results