

# ISM6251: Machine Learning for Business Applications

Spring 2026 Syllabus

Dr. Tim Smith

## Course Information

Course Number: ISM6251 Credit Hours: 3

### AI-Assisted Learning Environment

This course uses **Google Colab with Gemini AI** as the primary development environment. Students will learn machine learning concepts by collaborating with AI to generate, understand, and refine code. This approach prepares you for the modern workplace where AI tools are integral to data science workflows.

### Cross-Listed Course Notice

This course is cross-listed as follows:

Course Number	Level	Program
<b>ISM6251</b>	Graduate	Machine Learning for Business Applications
<b>ISM4543</b>	Undergraduate	Applied Data Science

**Important for all students:** Both graduate and undergraduate students complete the **same assignments, quizzes, and exams**. All course materials, Canvas modules, and documentation refer to **ISM6251**. If you are enrolled in ISM4543, you are in the right place—simply follow all ISM6251 materials and instructions.

This is a challenging, technical course that requires students to be proficient in:

- **Python programming** for data handling and manipulation
- **Statistical concepts** including descriptive and inferential statistics
- **Data science fundamentals** including simple model building and evaluation
- **Mathematical foundations** for understanding machine learning algorithms

Students should have strong quantitative skills and be comfortable with programming. If you are not confident in these prerequisites, please consult with the instructor during the first week of class.

## Instructor Contact Information

**Instructor:** Dr. Tim Smith **Office:** School of Information Systems and Management (CIS Building) **Email:** smith515@usf.edu **Office Hours:** Mondays 4:30 to 6:00 pm (or by appointment)

**Teaching Assistant:** [TBD] **TA Email:** [TBD]

### ! Important

Email is the best way to contact me, and I will try to answer emails within 24 business hours. For all communications, you must include the course number and section in the title of your email. Also, clarify your question/request – vagueness will result in wasting time with back-and-forth emails. Not including such details will result in your email being ignored.

## Course Description

This course provides a hands-on introduction to machine learning techniques with a focus on practical business applications. Students will learn how to frame problems, prepare data, build and evaluate models, and communicate results to business stakeholders.

**What makes this course different:** You will learn ML concepts using **AI-assisted coding with Gemini** in Google Colab. Rather than memorizing syntax, you will focus on:

- Understanding when and why to use different ML techniques

- Effectively prompting AI to generate correct code
- Verifying and interpreting AI-generated solutions
- Communicating ML insights to business stakeholders

This approach mirrors how data scientists work in industry today, where AI tools augment human expertise.

## Course Learning Outcomes

By the end of this course, students will be able to:

- Understand key machine learning concepts and their relevance to business decision-making
- **Understand how LLMs and transformers work** and their limitations as coding assistants
- **Effectively prompt AI tools** to generate ML code for regression and classification
- Prepare and clean data using appropriate statistical and programmatic techniques
- **Apply feature engineering strategies** to create business-relevant model inputs
- Evaluate and compare models using appropriate performance metrics
- **Verify and validate** AI-generated code and model outputs
- Identify business risks associated with overfitting, class imbalance, and interpretability
- Present and communicate machine learning insights to both technical and non-technical audiences
- Implement unsupervised learning approaches including **clustering and dimensionality reduction (PCA)**
- **Understand MLOps concepts** including model deployment, monitoring, and maintenance
- **Document and explain** AI-assisted development workflows

## Prerequisites

Students are expected to have completed a course in Data Mining or have equivalent experience. Proficiency in Python programming is required, as Python will be the primary language used in this course. Additionally, familiarity with statistics, linear algebra, and basic data preprocessing techniques is recommended.

# Technology Requirements

## Required Tools

1. **Google Account** - Required for Google Colab access (use your USF email)
2. **Google Colab** - Cloud-based notebook environment (free, no installation needed)
3. **Gemini AI** - Built into Google Colab for AI-assisted coding
4. **Web Browser** - Chrome recommended for best Colab compatibility

## Hardware Requirements

- Laptop with internet connectivity (required for every class)
- Minimum 8GB RAM recommended
- Webcam and microphone for proctored assessments

### **i** Note

Unlike previous semesters, you do **not** need to install Python locally. All coding will be done in Google Colab, which runs in your browser.

## Course Policies

### First-Day Attendance Policy

Per university policy, a student who is absent on the first day of class will automatically be dropped. Attendance will be taken at the start of class on the first day of class.

### Attending the Class Virtually or Physically

Note that the schedule indicates if the class for that week is F2F (face-to-face) or AO (asynchronous online). AO classes are pre-recorded, and attendance is virtual. F2F classes will be live classes that you attend.

### **i** Note

The selective use of AO classes is to enhance student learning – as there are a small number of course topics that are more easily digested via a pre-recorded tutorial-style video, which allows the students to pause, repeat, and engage in complex material in a way that is more easily digested.

## Course Materials

Students do not need to purchase any software or book for this class. All required materials are provided through Canvas and Google Colab.

## DataCamp Integration

**Important:** The first three weeks of this course are dedicated to developing Python and data management/preparation skills through DataCamp courses. Students are required to:

- Sign up for DataCamp using their **USF email address**
- Access DataCamp through the invitation link provided in Canvas
- Complete the assigned DataCamp course for each week as a graded deliverable

## Required DataCamp Courses (Weeks 1-3):

- **Week 1:** “Introduction to Python” - Foundation for ML programming
- **Week 2:** “Intermediate Python” - Advanced Python concepts for data science
- **Week 3:** “Data Manipulation with Pandas” - Essential data preparation skills

Through special arrangements with DataCamp, students will have **full access to the entire DataCamp platform throughout the semester**. Students are encouraged to explore additional courses related to our course content.

## How to Succeed in This Course

- Access course materials regularly
- Watch all recorded videos
- **Practice the AI-assisted workflow** – this is critical!
- **Understand concepts, not just code** – you must explain your work
- Document your learning process
- Work effectively in a team environment for the final project

## AI-Assisted Learning Policy

### ! AI Usage Guidelines

This course embraces AI as a learning tool. However, there are specific guidelines for when and how AI may be used.

## **During Assignments**

**AI assistance is REQUIRED and EXPECTED:**

- Use Gemini in Google Colab to generate code
- Ask AI to explain concepts you don't understand
- Iterate with AI to refine solutions
- Export your Gemini conversation as a submission deliverable

**Your responsibilities:**

- Understand all code before running it
- Verify that outputs are correct
- Be able to explain any part of your submission
- Complete the reflection section honestly

## **During Quizzes and Exams**

**AI assistance is NOT PERMITTED:**

- Quizzes and exams are closed-book, closed-AI
- Proctored with Honorlock
- Tests your understanding of concepts learned with AI assistance
- You should be able to answer questions about concepts you learned in assignments

## **Reflection & AI Documentation Requirement**

Every assignment requires three deliverables: the notebook, your exported Gemini conversation script, and a video walkthrough. The Gemini transcript and video document your AI interaction process.

Each notebook ends with a **reflection section** where you write in your own words about what was difficult, what was easier than expected, what surprised you, and what you would do differently next time.

Incomplete or dishonest reflections will result in significant point deductions.

## **Weekly Schedule**

**Mode:** F2F = Face-to-Face | AO = Asynchronous Online

## Foundation Weeks (1-3)

### Week 1

---

<b>Topic</b>	Course Introduction & AI-Assisted ML
<b>Date</b>	Jan 13 (F2F)
<b>Content</b>	Course overview • Introduction to ML & Predictive Modeling • How LLMs and Transformers Work • Getting started with Gemini in Colab • The AI-assisted coding workflow
<b>Deliverables</b>	Setup Google Colab & Gemini • Topic Practice • DataCamp: “Introduction to Python” • Quiz

---

### Week 2

---

<b>Topic</b>	Python Fundamentals with AI Assistance
<b>Date</b>	Jan 20 (AO)
<b>Content</b>	Python essentials with Gemini • NumPy, Pandas, Matplotlib • Learning to verify AI-generated code
<b>Deliverables</b>	Topic Practice (Python, NumPy, Pandas, Matplotlib) • DataCamp: “Intermediate Python” • Quiz

---

### Week 3

---

<b>Topic</b>	Data Preparation & Feature Engineering
<b>Date</b>	Jan 27 (AO)
<b>Content</b>	Handling missing data and imputation • Using AI for data cleaning • Feature Engineering Strategies • Feature Selection Methods
<b>Deliverables</b>	Assignment • Assignment Sample • Topic Practice • DataCamp: “Data Manipulation with Pandas” • Quiz

---

## Supervised Learning Weeks (4-10)

### Week 4

---

<b>Topic</b>	Linear Regression
<b>Date</b>	Feb 3 (F2F)
<b>Content</b>	Regression concepts and evaluation • MSE, RMSE, MAE, $R^2$ • Outlier treatment and feature selection • AI-assisted model building
<b>Deliverables</b>	Assignment • Assignment Sample • Topic Practice • Quiz

---

### Week 5

---

<b>Topic</b>	Logistic Regression
<b>Date</b>	Feb 10 (F2F)
<b>Content</b>	Binary classification fundamentals • Confusion Matrix, Precision, Recall, F1 • Encoding categorical variables • Threshold optimization for business
<b>Deliverables</b>	Assignment • Assignment Sample • Topic Practice • Quiz

---

### Week 6

---

<b>Topic</b>	K-Nearest Neighbors (KNN)
<b>Date</b>	Feb 17 (F2F)
<b>Content</b>	Distance-based classification • Feature scaling importance • Choosing K and distance metrics • Multi-class classification
<b>Deliverables</b>	Assignment • Assignment Sample • Topic Practice • Quiz

---

### Week 7

---

<b>Topic</b>	Support Vector Machines (SVM)
--------------	-------------------------------

---

---

<b>Date</b>	Feb 24 (F2F)
<b>Content</b>	Maximum margin classifiers • Kernel methods • Class imbalance handling • Precision-Recall curves
<b>Deliverables</b>	Assignment • Assignment Sample • Topic Practice • Quiz

---

### Week 8

---

<b>Topic</b>	Decision Trees
<b>Date</b>	Mar 3 (F2F)
<b>Content</b>	Tree-based models and interpretability • ROC curves and AUC • Hyperparameter tuning • Grid search and cross-validation
<b>Deliverables</b>	Assignment • Assignment Sample • Topic Practice • Quiz

---

### Week 9

---

<b>Topic</b>	Ensemble Methods
<b>Date</b>	Mar 10 (F2F)
<b>Content</b>	Random Forests (bagging) • Gradient Boosting (boosting) • Feature importance analysis • When to use ensembles
<b>Deliverables</b>	Assignment • Assignment Sample • Topic Practice • Quiz

---

### Week 10

---

<b>Topic</b>	Neural Networks
<b>Date</b>	Mar 17 (F2F)
<b>Content</b>	MLP fundamentals • Architecture decisions • Trade-offs in model complexity • MLOps Concepts & Model Lifecycle
<b>Deliverables</b>	Assignment • Assignment Sample • Topic Practice • Quiz

---

## Unsupervised Learning & Wrap-Up (11-12)

### Week 11

---

Topic	Clustering
Date	Mar 24 (F2F)
Content	K-Means and choosing K (elbow, silhouette) • Feature scaling for clustering • Cluster evaluation and interpretation • Customer segmentation applications
Deliverables	Assignment • Assignment Sample • Topic Practice • Quiz • Final Project work

---

### Week 12

---

Topic	Final Exam & Project Due
Date	Mar 31 (F2F)
Content	Comprehensive assessment of all ML concepts • Closed book, closed AI (Honorlock proctored)
Deliverables	Final Exam • Final Project Due

---

 Warning

**NOTE:** Course Schedule and Syllabus is Subject to Change at the Discretion of the Professor

## Grading

### Grading Components

---

Component	Weight	Notes
Weekly Assignments	15%	Business case studies with AI assistance
In-Class Quizzes	15%	Tests conceptual understanding

Component	Weight	Notes
Final Project	30%	Group project with AI documentation
Final Exam	30%	Closed-book, closed-AI
Discussion Participation	5%	Focus on AI-assisted learning experiences
DataCamp Courses	5%	Weeks 1-3 required courses

### Assignment Grading Breakdown

Each weekly assignment is graded using this rubric:

Component	Weight	Description
Problem Framing	10%	Articulate the business problem, justify the ML approach
Technical Implementation	25%	Working code, correct methodology
Verification Trail	15%	Sanity checks, validation of results
Understanding	25%	Explanations in YOUR OWN WORDS
AI Interaction	10%	Documented prompt evolution
Critical Analysis	15%	Error identification, limitations awareness

**Plus:** Video Walkthrough (3-7 minutes) is required for each assignment.

### Grading Scale

Score	Grade
97	A+
94	A
90	A-
87	B+
84	B
80	B-
77	C+
74	C
70	C-

Score	Grade
67	D+
64	D
60	D-
Below 60%	F

**i** Note

Please note that this tentative grading scale can change based on the students and overall class performance. The faculty reserves the right to change the grading scale as he deems appropriate based on the class's overall performance.

## Assessment and Participation Policies

### Assignments

Each week (starting Week 3) includes three learning materials:

Material	Purpose
<b>Assignment</b>	Graded business case study you complete and submit
<b>Assignment Sample</b>	Worked example with different company (study this first)
<b>Topic Practice</b>	Practice problems focused on the ML technique

### Assignment requirements:

- Completed in **Google Colab with Gemini assistance**
- **Gemini conversation export is required** (submitted as a deliverable)
- **Video walkthrough is required** (3-7 minutes explaining your work)
- Must demonstrate understanding, not just correct outputs
- Late submissions receive no credit

### Quizzes

- Quizzes will be conducted **in-class using Honorlock**
- Formats include multiple choice, multiple select, fill-in-the-blank, and short answer

- Quizzes are **closed book, closed AI, closed internet**
- Tests your understanding of concepts learned through AI-assisted practice
- Students are allowed to **drop one lowest quiz score**

## Final Exam

- Administered **in-class on the last day of class**
- **Closed book, closed AI, closed internet** using **Honorlock**
- Comprehensive coverage of all ML concepts
- You should be able to answer questions about concepts you practiced with AI

## Final Group Project

The final project will be a group project (3-4 students). You will:

- Identify a publicly available dataset
- Perform analysis, visualization, and model creation using AI assistance
- Document your AI workflow and collaboration process
- Receive peer evaluations from team members

Your final grade is the team's mark multiplied by your peer evaluation percentage.

## Technology and Equipment

- Students must bring a **fully charged laptop to every class**
- Ensure stable internet connectivity for Colab access
- **Power outlets are limited** – plan accordingly

## Course Policies

### Late Work Policy

#### ! Important

No credit will be given to late submissions. Plan your time well and target to complete submissions at least one day before the deadline.

## **Grades of “Incomplete”**

An “I” grade may be awarded when:

1. Arrangements are made before the end of the semester
2. In the judgment of the instructor, a valid reason is offered
3. A clear path to completion is agreed upon

## **Email**

The primary means of communication is email. Include the course number and section in all email subject lines. Expect 24 business hour response times.

## **Canvas**

This course uses Canvas for materials, quizzes, and submissions. If you need help, consult Canvas help guides or contact USF IT at (813) 974-1222.

## **Academic Integrity**

Academic integrity is fundamental to this course. While AI assistance is expected for assignments, you must:

- Understand and be able to explain all submitted work
- Submit Gemini conversation transcripts and complete reflections honestly
- Not use AI during proctored quizzes and exams
- Submit original explanations and interpretations

Violations will result in a zero for the assignment and may be reported to the Office of Student Conduct.

## **Disruption to Academic Process**

Disruptive behavior in the academic setting hinders the educational process and will not be tolerated.

## Standard University Policies

Policies about disability access, religious observances, academic grievances, academic integrity and misconduct, academic continuity, food insecurity, and sexual harassment are governed by a central set of policies that apply to all classes at USF. These may be accessed at: <https://www.usf.edu/provost/faculty/core-syllabus-policy-statements.aspx>

## Online Proctoring

Quizzes and exams will be conducted using Honorlock. Students must have a webcam and microphone. Setup information will be provided before proctored assessments.

For additional information about online proctoring: <https://www.usf.edu/innovative-education/student-resources/index.aspx>

## Course Hero / Chegg Policy

The USF Policy on Academic Integrity specifies that students may not use websites that enable cheating, such as Chegg.com and CourseHero.com.

## End of Semester Student Evaluations

All classes at USF use an online system for student feedback. Your participation is highly encouraged and valued.

## Getting Help

### Technical Issues

- **Google Colab:** Check [Colab FAQ](#)
- **Canvas:** Contact USF IT at (813) 974-1222 or [help@usf.edu](mailto:help@usf.edu)
- **Honorlock:** Check the online proctoring student FAQ

### Course Content

- Review course materials and AI-assisted notebooks
- Ask Gemini to explain concepts
- Post on Canvas Discussions
- Attend office hours

## **Academic Support**

- USF Office of Student Success coordinates university-wide academic support
- Visit the Office of Student Success website for resources

---

*Syllabus Version: 2.0 Last Updated: February 2026*

Any USF policy will supersede any policy mentioned in this syllabus.