

Course Syllabus

Day 1: The Mechanics of the Machine

Focus: Demystifying AI, understanding the core algorithms, and setting up the analytical environment in Python.

Time	Module	Core Concepts & Activities
0800 - 0930	The Analytics Pyramid & Supervised Learning	Defining AI vs. Machine Learning vs. Deep Learning. We explore how supervised learning models are algorithmically trained using historically labeled military datasets. The Math: Understanding how models learn the mapping function from operational inputs to tactical outputs.
0945 - 1130	Unsupervised Learning: Processing the Unknown	In combat, you can't map every variable in advance. We discuss how unsupervised models train on raw data, identifying patterns without human-provided answer keys. Practical App: Using Primary Dimensionality Reduction (PCA) to strip atmospheric noise from hyperspectral imagery.
1130 - 1230	Lunch	

Time	Module	Core Concepts & Activities
1230 - 1430	Deep Learning & Artificial Neural Networks (ANNs)	Breaking down the Perceptron and how neural networks use activation functions to transform outputs. We will explore how deep learning models pass the outputs of one hidden layer as the inputs to additional hidden layers.
1445 - 1630	Project Kickoff: Setting the Baseline in Python	<i>Operation Red Pill</i> begins. Students configure their analytical environments, load their initial datasets into Python, and establish baseline performance metrics.

Day 2: The Tactical Edge & The Foundation Era

Focus: Applying AI to military logistics and intelligence, and navigating the complexities of large-scale models.

Time	Module	Core Concepts & Activities
0800 - 1000	Foundation Models & Natural Language Processing (NLP)	A foundation model is a large ML model trained on a vast quantity of data at scale so that it can be adapted to a wide range of downstream tasks. Exploring NLP concepts: variability (many forms mean the same thing) and

Time	Module	Core Concepts & Activities
		ambiguity (one form can mean many things).
1015 - 1130	Military Applications in the Trenches	From zero-shot learning to applying AI for ISR target detection and predictive logistics. We look at how regression algorithms can predict when an aircraft or vehicle component will fail, optimizing proactive maintenance.
1130 - 1230	Lunch	
1230 - 1430	Competing Interests of Machine Learning	A great model isn't just accurate. We must balance Accuracy, Privacy & Security, Robustness, Interpretability, and Fairness. The Danger Zone: Understanding vulnerabilities like Model Inversion, Extraction, and Adversarial Patches.
1445 - 1630	Project Execution: Poisoning the Well	Students actively introduce anomalies and data poisoning into their Python models. They will document the degradation of their model's performance and track how factuality hallucinations (factual inconsistency and factual fabrication) occur.

Day 3: The Feynman Protocols: Is AI behaving the way we want it to

Focus: "What I cannot create, I do not understand." Testing, breaking, and explaining AI to leadership.

Time	Module	Core Concepts & Activities
0800 - 0930	Feynman Questioning: Is it Actually Working?	Teaching students to stop trusting the dashboard and start looking at the track. How to design a scientific experiment to test if an AI model is generalizing or merely overfitting to a specific training environment.
0945 - 1130	Evaluating Faithfulness & Logical Consistency	Assessing when LLM outputs deviate from user directives (Instruction inconsistency) or exhibit internal logical contradictions (Logical inconsistency).
1130 - 1230	Lunch	
1230 - 1430	Project Finalization: Translating Math to Mission	Students finalize their Python experiments. They must consolidate their findings into a brief that strips away the technical jargon, focusing entirely on operational impact, risk, and deployment readiness.
1445 - 1630	Leadership Briefings &	Each cell briefs the

Time	Module	Core Concepts & Activities
	Course Hotwash	instructor (acting as the Commander). Feedback will center on clarity, candor, and the ability to explain <i>why</i> the AI failed and <i>what</i> constraints must be respected before deployment.