Introduction and Motivation

AERO 689: Introduction to Machine Learning for Aerospace Engineers

Dr. Raktim Bhattacharya

Texas A and M University - Aerospace Engineering

2026-01-01

Welcome to AERO 689

The Intersection of Data Science and Aerospace Engineering

Where machine learning meets flight dynamics, aerodynamics, and space systems

Course Logistics & Resources

How to Succeed:

- Participate in class
- Practice coding
- Ask questions early
- Use office hours

Assessment Strategy

- Homework: 30%
- Paper review: 10%
- Midterm: 30%
- Final Project: 30%

Resources:

- Canvas
- GitHub
- NASA, AIAA, and textbook links

Why This Course?

Modern aerospace systems generate **large volumes of data**:

- Flight-test data
- Wind-tunnel experiments
- CFD and simulation databases
- Telemetry and health-monitoring data

Traditional physics-based models alone are often:

- Expensive
- Incomplete
- Uncertain
  Machine learning provides tools to extract structure from data.

Where Machine Learning Fits in Aerospace

Machine learning is used as:

- **Surrogate modeling**
  (fast approximation of expensive simulations)
- **Residual modeling**
  (physics model + data-driven correction)
- **Pattern recognition**
  (classification, clustering, anomaly detection)

*It does **not replace physics**. It complements it.*

Learning Objectives (Part 1)

LO 1: Understand ML in Aerospace

- Explain the fundamental concepts of machine learning and how they differ from traditional programming in aerospace contexts.
- Describe the historical evolution and motivation for ML in aerospace engineering.
- Recognize the benefits and limitations of ML for safety-critical aerospace systems.
- Discuss real-world examples where ML has impacted aerospace (e.g., anomaly detection, predictive maintenance, autonomous operations).

Learning Objectives (Part 2)

LO 2: Identify Key Applications

- Identify major aerospace domains where ML is applied, such as flight dynamics, aerodynamics, navigation, and space systems.
- Match ML techniques (supervised, unsupervised, reinforcement) to relevant aerospace problems.
- Summarize recent success stories and industry trends in aerospace ML.

Learning Objectives (Part 3)

LO 3: Recognize Unique Challenges

- Understand the unique challenges of aerospace data: sparsity, noise, high dimensionality, and safety requirements.
- Explain the importance of physics-informed ML and regulatory considerations in aerospace.
- Appreciate the ethical responsibilities and safety implications of deploying ML in aerospace systems.

The Evolution of Automation in Aerospace

- 1960s: Early autopilots and control theory
- 1970s: Kalman filter revolutionizes navigation
- 1980s: NASA experiments with neural networks for fault detection
- 2000s: Data-driven methods in engine health monitoring
- 2010s–2020s: Deep learning, real-time ML, and autonomous systems (SpaceX, Boeing, Airbus)

What Is Machine Learning?

Working definition:

Machine learning is the study of algorithms that learn mappings from data by solving optimization problems.

Key idea:

- Learn a function
- Evaluate its performance
- Ensure it generalizes beyond training data

What is Machine Learning? (contd.)

Traditional Programming:

- Input (Data) + Program (Rules) $\rightarrow$ Output (Prediction)
- E.g.
  - Input: 1,2,3,4, ...
  - Program: $f(x) = 2x + 1$ We code this
  - Output: 3,5,7,9, ...

Machine Learning:

- Input (Data) + Output (Labels) $\rightarrow$ Program (Model)
- E.g.
  - Input: 1,2,3,4, ...
  - Output: 3,5,7,9, ...
  - Program: ML learns $f(x) = 2x + 1$ from the data

A Generalized Function Approximation

- ML can be treated as a function approximation technique
- Except input/output data can be of any data type
  - Images
  - Sounds
  - etc.
- ML learns the relationship between input and output data
  - can be used to predict/generate outputs from other inputs
  - we may not capture the relation exactly
    - prediction will have errors!

What Machine Learning Is NOT

Machine learning is not:

- Magic
- A substitute for understanding physics
- Guaranteed to extrapolate correctly
- Automatically safe or robust

In aerospace:

- Validation is critical
- Extrapolation is dangerous
- Guarantees matter

Types of Machine Learning

Supervised Learning: Learn from labeled data

- Regression (continuous outputs)
- Classification (discrete labels)

Unsupervised Learning: Find patterns in unlabeled data

- Clustering
- Dimensionality reduction

Reinforcement Learning: Learn by trial and error

- Sequential decision making

Supervised Learning (Examples)

Labeled Inputs $\rightarrow$ Outputs

Examples:

- (Geometry, Mach number) $\rightarrow$ Lift coefficient
- Sensor measurements $\rightarrow$ Flight regime
- Telemetry $\rightarrow$ Fault / no-fault

You define:

- Inputs
- Outputs
- Performance metric

Unsupervised Learning (Examples)

No labeled outputs

Examples:

- Grouping flight trajectories
- Identifying maneuver types
- Discovering dominant flow structures

Goal:

- Find structure in data
- Reduce dimensionality

Reinforcement Learning (Examples)

Learns via Interaction

Basic Idea:

- Agent interacts with a system
- Receives rewards
- Learns a policy – like a controller or guidance algorithm

Examples:

- Autopilot tuning
- Trajectory optimization

The Machine Learning Pipeline

A typical workflow:

1. Problem formulation
2. Data collection
3. Preprocessing
4. Model selection
5. Training
6. Validation and testing
7. Interpretation and deployment

Most failures occur at:

- Step 1 (wrong question)
- Step 6 (poor validation)

Why is ML Hard in Aerospace?

- Data is expensive and sparse (flight tests, wind tunnels)
- Safety-critical: mistakes can cost lives
- Physics-based constraints: not all patterns are physically meaningful
- Certification and regulation (FAA, NASA)
- Real-time and embedded requirements

Why Validation Matters in Aerospace

Aerospace systems are:

- Safety-critical
- Operated outside training data
- Sensitive to uncertainty

Key questions:

- Does the model generalize?
- How sensitive is it to noise?
- What happens outside the data range?

A hybrid approach of physics + data driven machine learning is likely to work better for aerospace systems.

Ethics and Responsibility in Aerospace ML

Machine learning can influence:

- Control decisions
- Navigation and guidance
- Fault detection and isolation
- Autonomous behavior

Errors can cost lives, assets, or strategic stability

- Would you trust an ML autopilot? Why or why not?
- How do we ensure safety, transparency, and accountability?
- What are the risks of bias or overfitting in safety-critical systems?

What Makes Aerospace ML Different (from consumer ML)?

- Smaller datasets
- Rare but catastrophic failure modes
- Operation outside training conditions

Ethics and Responsibility in Aerospace ML (contd.)

Core Ethical Questions

When deploying ML in aerospace, ask:

- Can this system fail safely?
- Do we understand when it will fail?
- Who is accountable for decisions?
- Can humans intervene?
- Are the assumptions documented?

Data Ethics

- Data quality and representativeness
- Bias in training data
- Missing or censored failure cases
- Simulation vs real-world mismatch

A model trained on biased data $\rightarrow$ produces biased decisions

Ethics: Examples and Interpretability

Example: Biased Flight Data

If training data includes:

- Only nominal flight conditions
- Few extreme maneuvers
- Limited environmental variation

Then the model may:

- Perform well in tests
- Fail catastrophically in edge cases

Model Transparency and Interpretability

Important questions:

- Can we explain model outputs?
- Do engineers understand failure modes?
- Is debugging possible?

Black-box models:

Verification and Validation (V&V)

Traditional aerospace V&V

- Analytical guarantees
- Worst-case analysis
- Structured testing

ML challenges

- Statistical performance $\neq$ safety guarantee
- No universal stability guarantees

Validation Is Not Optional

High test accuracy does NOT imply:

- Robustness
- Safety
- Reliability

Ethical deployment requires:

- Stress testing
- Uncertainty analysis
- Out-of-distribution evaluation

Human-in-the-Loop Responsibility

Critical design choice:

- Fully autonomous?
- Human-supervised?
- Human override capability?

Ethical principle:

ML should assist decision-making, not obscure it

Automation Bias

Risk:

- Humans over-trust ML outputs
- Operators defer judgment to the algorithm

Mitigation:

- Clear confidence measures
- Training users to question outputs
- Explicit failure indicators

Fault Detection and False Confidence

In aerospace:

- False negatives can be catastrophic
- False positives can cause mission aborts

Ethical tradeoff:

- Conservative models vs operational efficiency
- Transparency about detection limits

Responsibility and Accountability

When ML fails, who is responsible?

- The algorithm designer?
- The system integrator?
- The operator?
- The organization?

Ethical engineering requires:

- Clear responsibility chains
- Documented assumptions

ML in Autonomous Systems

As machine learning becomes more deeply integrated into autonomous aerospace systems, several unique challenges and risks must be considered:

Escalation of autonomy:

- Systems may transition from human-supervised to fully autonomous operation, sometimes without clear boundaries or sufficient oversight.
- Increased autonomy can outpace regulatory frameworks and operator understanding, raising questions about control, intervention, and accountability.
- Example: An autopilot that adapts its own logic in flight, making decisions beyond its original certification.

ML in Autonomous Systems (contd.)

Unexpected emergent behavior:

- Complex ML models can exhibit behaviors not anticipated during design or testing, especially when exposed to novel situations or data distributions.
- These emergent behaviors may be beneficial, neutral, or hazardous, and are often difficult to predict or diagnose.
- Example: Anomaly detection system that misclassifies rare but critical events, or a control system that exploits loopholes in its reward function.

ML in Autonomous Systems (contd.)

Limited real-time interpretability:

- Many ML models, especially deep neural networks, act as "black boxes," making it hard for engineers and operators to understand or trust their decisions in real time.
- Lack of interpretability can hinder rapid troubleshooting, regulatory approval, and safe human-machine teaming.
- Example: A flight control system issues a corrective action, but the rationale is opaque to the pilot and ground crew.
  Design principle: Autonomy must degrade gracefully

The Hype Around ML/AI

Why is everyone talking about AI?

- Headlines: "AI will revolutionize everything!"
  - AI is the "magic wand" that will solve everything
- Startups and big tech racing to deploy ML/AI
- Media focus on ChatGPT, self-driving cars, and robots
- Massive investments and job growth in AI/ML
- Fear of "AI replacing humans" and ethical debates

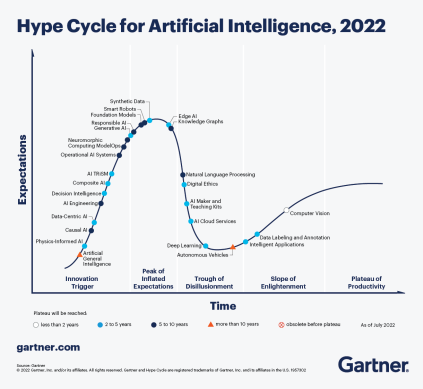

Figure 1: AI Hype Cycle

The Hype Around ML/AI (contd.)

Reality Check

- Not every problem needs ML/AI
- Many "AI" products are just automation or statistics
- Hype can lead to unrealistic expectations and failures
- Responsible engineers must separate fact from fiction

Potential Benefits of AI

- Automates repetitive and complex tasks
- Enables data-driven decision making
- Improves accuracy and efficiency
- Unlocks new capabilities (e.g., perception, prediction)
- Accelerates scientific discovery and innovation

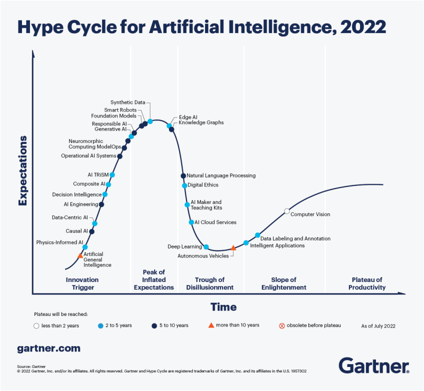

Figure 2: AI Hype Cycle

Interactive Question

What aerospace systems around you could benefit from machine learning?

Flight Systems

- Autopilot
- Navigation
- Flight control

Maintenance

- Engine health
- Predictive maintenance
- Structural monitoring

Environment

- Weather prediction
- Turbulence detection
- Ice detection

Operations

- Air traffic control
- Route optimization
- Fuel management

What Makes Aerospace Different?

Unique Challenges

- Safety-Critical – Lives depend on our algorithms
- Physics-Informed - Centuries of knowledge
  - fluid mechanics (e.g., Navier-Stokes equations for airflow)
  - analytical dynamics (e.g., equations of motion for aircraft)
  - material science (e.g., fatigue and failure models)
  - flight stability theory, orbital mechanics (e.g., satellite trajectories)
  - propulsion system modeling, etc.
- Multi-Scale – Molecular flows to orbital mechanics
- Regulatory – FAA/NASA certification
- Real-Time – Split-second decisions required

What Makes Aerospace Different? (contd.)

Aerospace Data Characteristics

- **Time-series**: Continuous sensor streams (100+ Hz)
- **Multi-dimensional**: Pressure, temperature, velocity, position, attitude
- **Noisy**: Sensor failures, environmental interference
- **Sparse**: Expensive flight tests, limited wind tunnel data

Best Practices for Aerospace ML

Best Practices for Responsible Aerospace ML

- Combine physics with learning
- Quantify uncertainty
- Validate beyond training data
- Keep humans informed
- Document assumptions
- Design for failure

What This Course Emphasizes

- Focus on interpretable models first
- Emphasize validation and uncertainty
- Avoid ML where physics is sufficient and tractable
- Treat ML as an engineering tool

Key Takeaways

Machine learning can save lives in aerospace - but only if applied with domain knowledge and rigorous validation

- **Safety First**: Aerospace ML must handle failure gracefully
  - Systems must be robust to faults and uncertainties
  - Safety and accountability always come first
- **Physics Informed**: Combine data with centuries of knowledge
  - Use physical laws (e.g., conservation of momentum, Navier-Stokes equations) to guide models
  - ML should complement, not replace, physics-based understanding
- **Responsibility**: Engineers have ethical obligations
  - Ethical responsibility is part of aerospace engineering
  - Validation and transparency are essential
  - ML is powerful but fragile – misuse can have serious consequences

Course Structure: Your Learning Journey

Phase 1: Foundations: Weeks 1-6

- Linear regression
- Classification
- Clustering & PCA
- Dimensionality reduction

Phase 2: Advanced: Weeks 7-12

- Neural networks
- Deep learning & CNNs
- Sequence models
- Physics-informed ML

Phase 3: Application: Weeks 13-14

- Project development
- Final presentations
- Real challenges