Intelligent Unmanned System Control Practice

Chapter 7 Health Management and Safety Assessment

Author: Xunhua Dai, Central South University

Website: https://faculty.csu.edu.cn/daixunhua

Email: dai.xh@csu.edu.cn

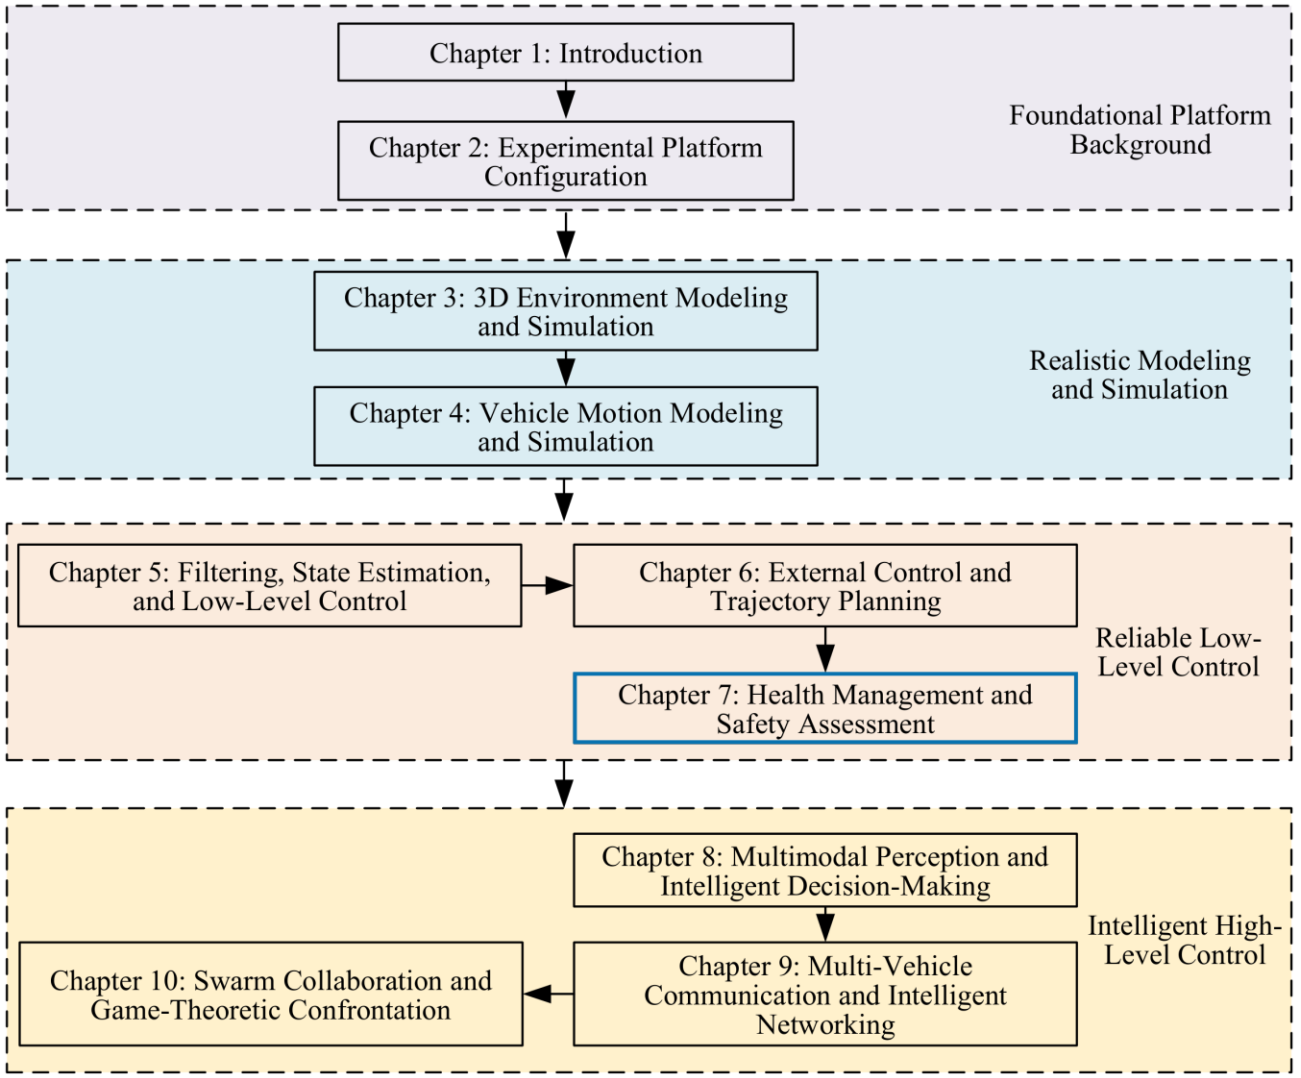
RflySim: http://rflysim.com

Companion Paper



Dai, Xunhua, Jinhu Tu, and Quan Quan. "RflySim ToolChain: a rapid development and validation toolchain for intelligent unmanned swarm systems." Journal of Systems Engineering and Electronics 36.4 (2025): 1077-1093.

Textbook is coming soon...



Chapter Learning Objectives

What You Can Do After Chapter 7

1. Distinguish anomalies, faults, and failures; relate health to mission risk.
2. Model faults; design traceable injection, data collection, and automated tests.
3. Assess diagnosis, prognostics, and risk from evidence; identify data and model limits.
4. Close the loop with safety decisions; test effects and costs at several levels.

Chapter focus: from normal operation to safe operation with faults.

Chapter Outline

Health Management and Safety Assessment: Learning Roadmap

Chapter 7 Health Management and Safety Assessment

7.1 Health Management Fundamentals

7.2 Unified Fault Modeling

7.3 Safety Reinforcement and Fault Data

7.4 Fault Hazards and Safety Assessment

7.5 Diagnosis and Prognostics

7.6 Safety Decisions and Improvement

7.7 Fault Injection Interfaces

7.8 Experiments and Validation

7.9 Exercises and Classroom Discussion

4 Units in This Section

7.1.1 Background and Significance

Textbook: Sec. 7.1.1

7.1.2 Health and Safety: Differences and Connections

Textbook: Sec. 7.1.2, 7.1.2.3

7.1.3 Health Management of Unmanned Systems

Textbook: Sec. 7.1.2.1

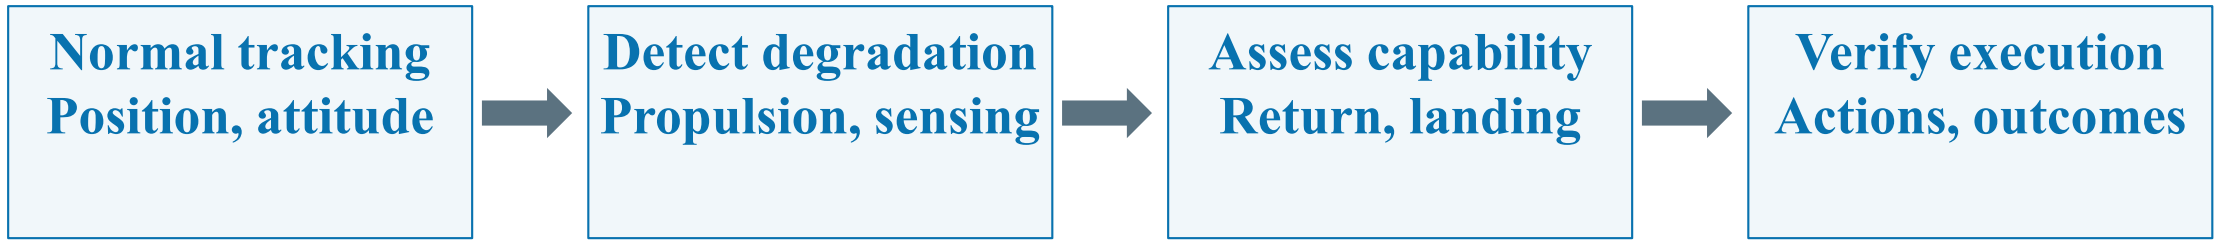
7.1.4 Algorithm Development Challenges and Validation

Textbook: Sec. 7.1.1

7.1.1 Background and Significance: Running Task

Control Keeps It Flying; Health Management Asks How Long

Running task: an inspection multicopter must complete or safely end its mission as propulsion degrades.



Established foundation: dynamics, state estimation, and control make motion controllable.

New question: after components change, can we still rely on control performance, control margin, and outcomes?

7.1.1 Background and Significance: Terminology

Anomalies, Internal Faults, and Functional Failures

Concept	Inspection mission example	Key distinction
Anomaly	Vibration or current exceeds the reference for the same conditions	Deviation from the reference under the same conditions
Degradation / fault	Bearing wear, reduced thrust efficiency, or sensor bias	Reduced performance / an abnormal internal state
Failure	Cannot maintain the required attitude or perform the required function	Loss of the ability to perform a required function
External disturbance	Gusts, collision contact, or external electromagnetic interference	External influence; a "fault" in this chapter's broad usage

An anomaly need not be a fault; no clear error does not rule out degradation.

7.1.1 Background and Significance: Propulsion Faults and Compensation

Why Can Flight Stay Stable Briefly after Thrust Falls?

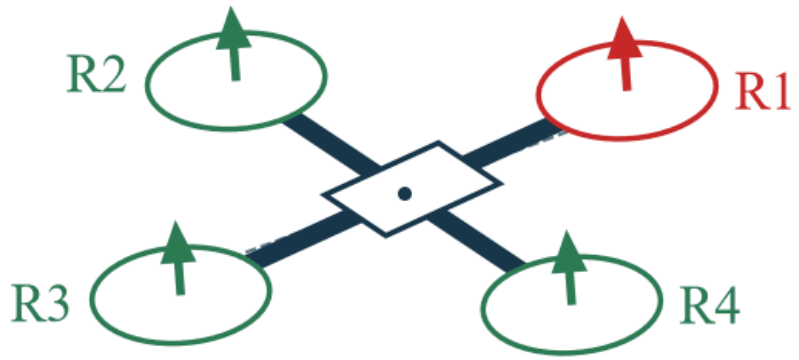
Rotor loss: compensation and saturation

$\eta_1 = 0.9$: hover restored; $\eta_1 = 0.6$: descent

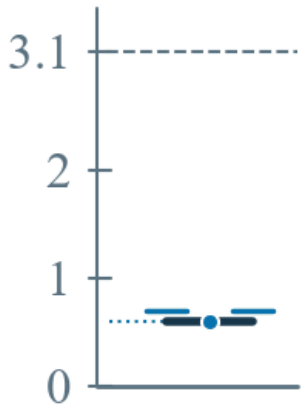
Actual thrust f_i

$\eta_1 = 0.6$

Level ref.



Command $u_1 = 1.50$ (≤ 1.50) Thrust $f_1 = 2.21$ N

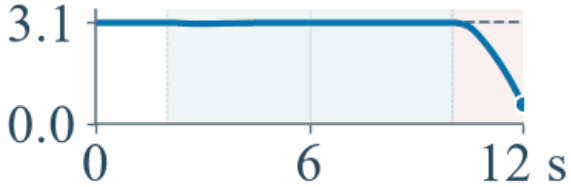


Mild-loss peaks

Tilt / ° 3.7

Drop / cm 3.5

Height z / m 0.60



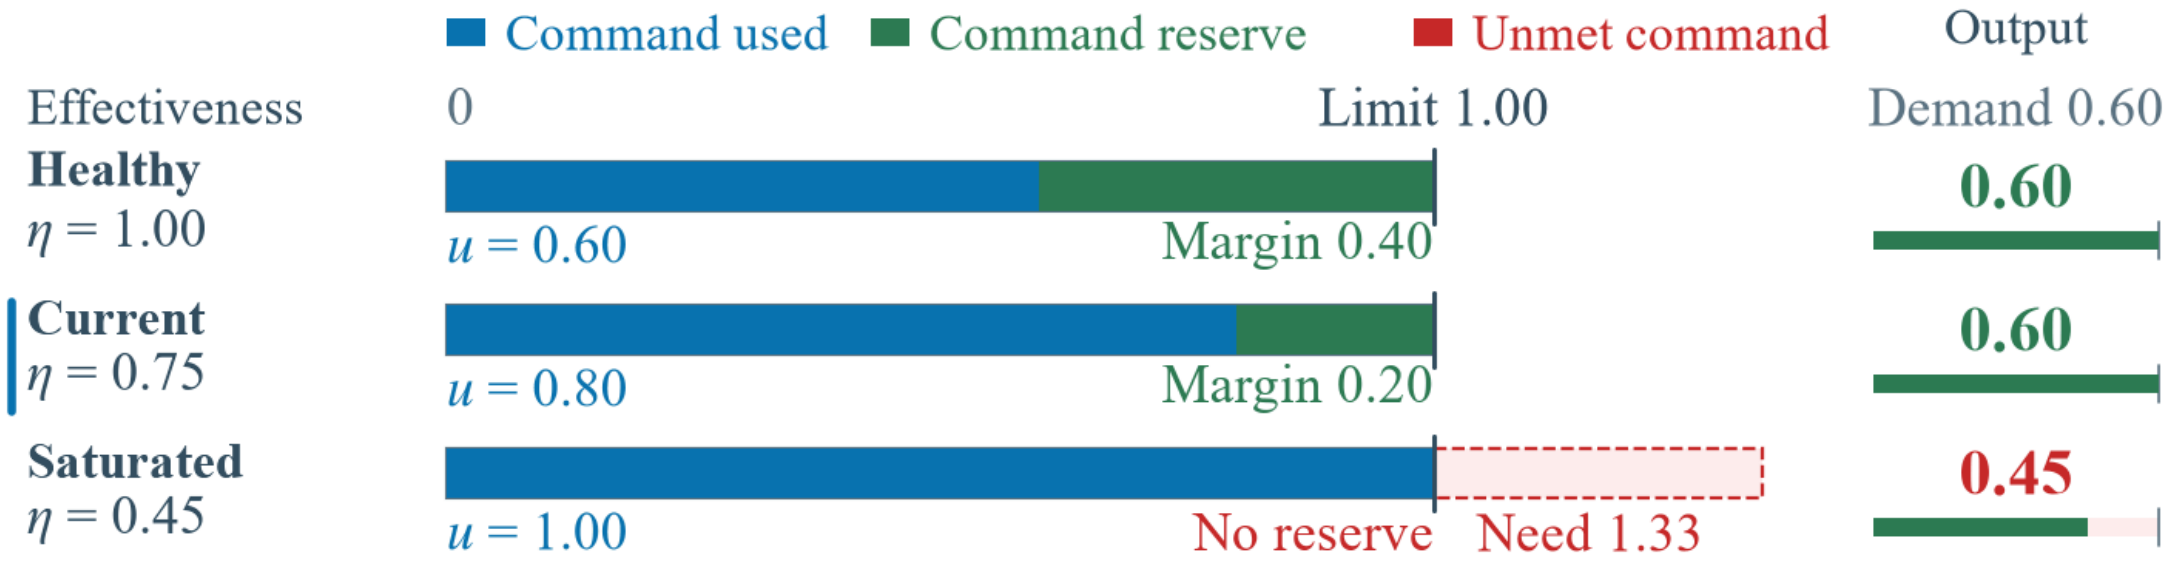
Fixed heading, torque priority; zero-torque thrust cap: 0.90 mg

Figure: one rotor loses effectiveness; closed-loop compensation (MuJoCo; effectiveness 1→0.9→0.6). Mild loss is compensated; severe loss cannot hold height with fixed heading and torque priority.

7.1.1 Background and Significance: Control Compensation

Control Margin Can Shrink while Error Stays Small

Compensation uses margin; saturation leaves an output shortfall



Output holds 0.60; margin 0.40 to 0.20. Saturation: output shortfall 0.15.

Figure: quasi-static thrust allocation. Effectiveness = output ÷ command; margin = command limit − required command.

7.1.1 Background and Significance: One-Step Calculation

Compensation and Limits at the Same Demand

$$y = \eta u, \quad 0 \le u \le 1, \quad u = \min(0.60/\eta, 1)$$

η: effectiveness (output ÷ command); u: normalized command (limit 1); y: normalized output.

State	Effectiveness	Required / actual command	Actual output
Healthy	1.00	0.60 / 0.60	0.60
Compensable loss	0.75	0.80 / 0.80	0.60
Beyond capability	0.45	1.33 / 1.00	0.45

Health monitoring: track results and the control margin spent to achieve them.

7.1.1 Background and Significance: Sensor Faults

Why Can Actual Height Be Wrong when the Reading Meets Its Target?

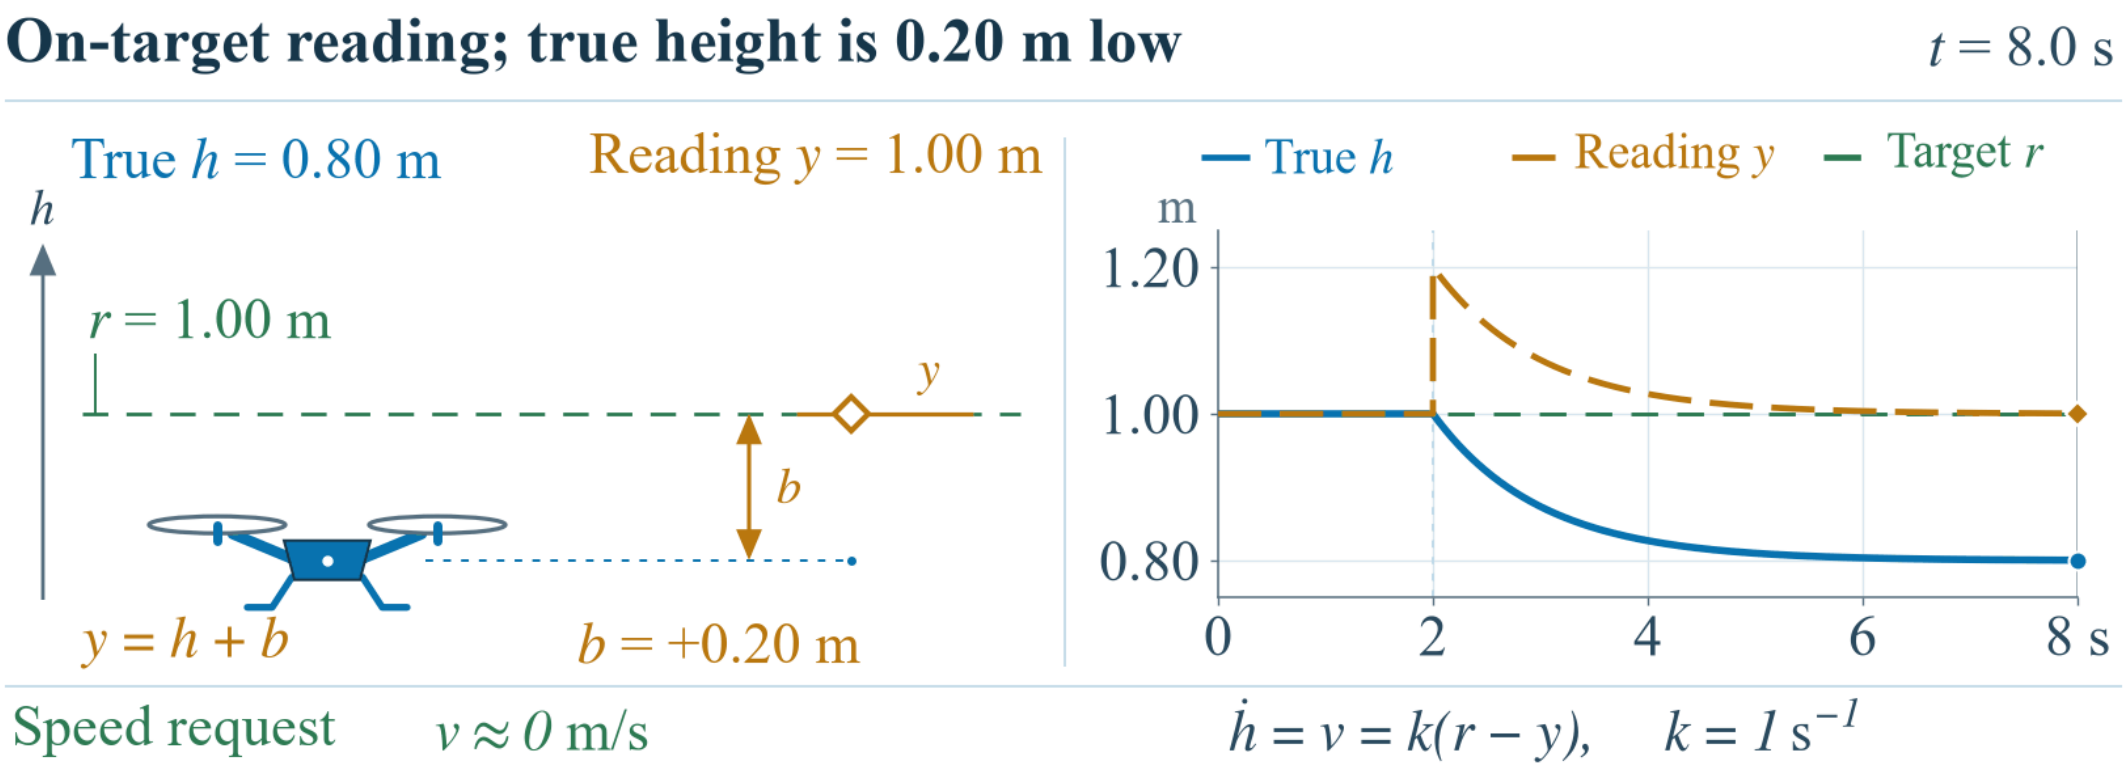


Figure: a persistent positive bias can make feedback move the vehicle away from the true target.

7.1.1 Background and Significance: Safety Requirements

Different Systems Have Different Paths to Harm

System	Representative hazards	Design and operating concerns
Embodied intelligent systems Arms, robot dogs	Collision, crushing, or uncontrolled gripping	Speed and force limits, human–robot separation, and controlled stopping
Autonomous road vehicles	Control loss; poor braking; missed detections	Space to stop, other road users, takeover, and degraded operation
Unmanned aerial vehicles (UAVs)	Attitude loss, insufficient thrust, falling	Remaining control capability; landing areas, people, and property below

UAVs: hard to stop on failure or add hardware redundancy; rely more on health management.

Safety depends on mission, environment, and safeguards, beyond system type.

Chapter Outline

Health Management and Safety Assessment: Learning Roadmap

Chapter 7 Health Management and Safety Assessment

7.1 Health Management Fundamentals

7.2 Unified Fault Modeling

7.3 Safety Reinforcement and Fault Data

7.4 Fault Hazards and Safety Assessment

7.5 Diagnosis and Prognostics

7.6 Safety Decisions and Improvement

7.7 Fault Injection Interfaces

7.8 Experiments and Validation

7.9 Exercises and Classroom Discussion

4 Units in This Section

7.1.1 Background and Significance

Textbook: Sec. 7.1.1

7.1.2 Health and Safety: Differences and Connections

Textbook: Sec. 7.1.2, 7.1.2.3

7.1.3 Health Management of Unmanned Systems

Textbook: Sec. 7.1.2.1

7.1.4 Algorithm Development Challenges and Validation

Textbook: Sec. 7.1.1

7.1.2 Health and Safety: Differences and Connections — Two Perspectives

Health Describes Condition; Safety Judges Consequences

Health Management	Safety Assessment
Focus: efficiency, wear, bias, vibration, and remaining capability.	**Focus:** loss-of-control paths, people/property at risk, response time, and safeguards.
Outputs: anomaly location, fault severity, trends, and uncertainty.	**Outputs:** risk judgments, constraints, and acceptable operating modes.

Textbook Table 7.1: data sources, assessment cycles, and key methods

Data: real-time onboard sensors	**Data:** history / external observations
Cycle: real-time / short-term monitoring	**Cycle:** design and operating stages
Tools: detection, diagnosis, life prediction	**Tools:** fault trees, failure mode analysis

A given health state may need different safety decisions as mission or environment changes.

7.1.2 Health and Safety: Differences and Connections — Goals and Data

Health Management Aids Recovery; Safety Assessment Limits Harm

Health Management

Goal: improve self-monitoring and recovery capability.

Applications: autopilot failsafe; onboard prognostics and health management (PHM).

Safety Assessment

Goal: reduce accident probability and the extent of the consequences.

Uses: airworthiness certification, mission risk assessment.

Data for the two RflySim modules (Textbook: Sec. 7.1.2.3)

Data: PX4 logs or state estimates
Nature: indirect, with error and noise
Use: real-time monitoring, early warning

Data: CopterSim / RflySim3D truth
Nature: omniscient; no sensor noise
Use: mission outcomes before/after faults

Health uses noisy estimates; safety assessment can check outcomes against simulation truth.

7.1.2 Health and Safety: Differences and Connections — Context

Return or Land Nearby? Context Matters for the Same Degradation

Nearby; clear alternate landing area	Far away, headwind, low battery	People below; high landing risk
Possible action	**Possible action**	**Possible action**
Return slowly if supported by evidence	Land nearby if margin is insufficient	Divert away from people, then select a landing site
Required evidence	**Required evidence**	**Required evidence**
Controllable attitude, sufficient margin, reachable landing site	Return energy, wind estimate, remaining capability along the route	Candidate sites, reachable area, people and property below

Assess: remaining capability → feasible actions → each action's consequences and constraints.

7.1.2 Health and Safety: Differences and Connections — Key Metrics

Reliability, Availability, Safety, and Dependability

Concept	Question	Example
Reliability	Can the required function be sustained over the specified period?	Does a functional failure occur during inspection?
Availability	Is the system ready for use when needed?	Waiting for repair reduces available time
Safety	Are risks from operation and fault consequences controlled?	A timely controlled landing after danger is detected
Dependability	What evidence supports relying on the system?	Evidence across reliability, safety, and other properties

7.1.2 Health and Safety: Differences and Connections — Safety and Security

An Attack Can Propagate through Control into Physical Harm

Safety	Addresses unintended harm from faults or abnormal operation; controls risks to people, equipment, and the environment.
Security	Defends against deliberate attacks/manipulation; protects observations, commands, and links from unauthorized interference.

A security failure may propagate through the control chain into a physical hazard:

Attack
Sensor spoofing

➡

Observation / link anomaly
Position distorted

➡

Wrong action
False error → turn

➡

Physical hazard
May cross boundary or collide

A hazard may develop only if abnormal data are trusted and safeguards fail.
Attack modeling: Section 7.2, “Sensor and Communication Faults.”

7.1.2 Health and Safety: Differences and Connections — Health and Life

Remaining Charge and Battery Aging Answer Different Questions

Quantity	Question answered	Additional conditions
State of charge (SOC)	How much usable charge remains this time?	Capacity estimate; charge / discharge state
State of health (SOH) (capacity here)	How much has it degraded relative to a reference battery?	Define capacity or internal resistance first
Remaining useful life (RUL)	How much use remains before end of life?	Threshold, future conditions, prediction interval

Predictive maintenance helps plan missions and maintenance ahead, beyond a life estimate.

7.1.2 Health and Safety: Differences and Connections — Remaining Life

Why Can the Same Health State Give Different Remaining Lives?

Same SOH, different conditional RUL

Assume constant future rates

Now: cycle 100

Scenario A
SOH = 0.90
Loss/cycle: 0.001

Scenario B
SOH = 0.90
Loss/cycle: 0.002

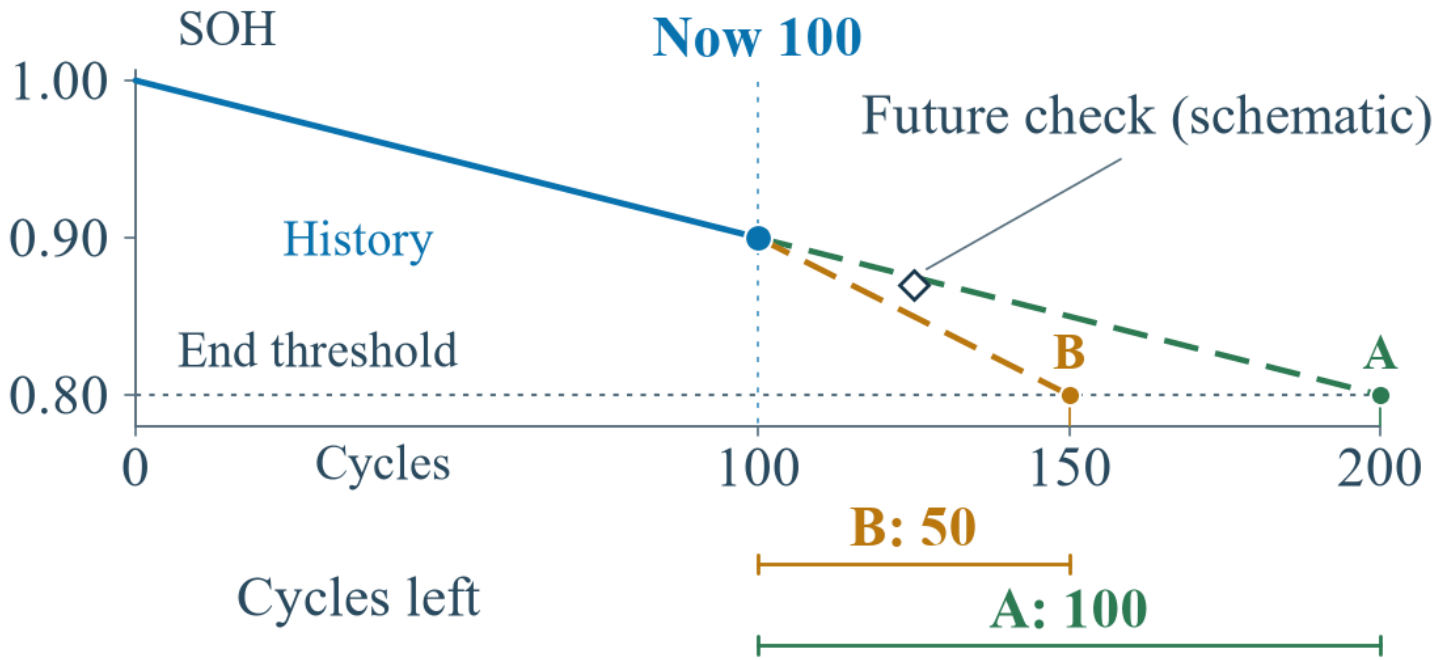


Figure: at the same current health, different future degradation conditions change the life prediction.

Chapter Outline

Health Management and Safety Assessment: Learning Roadmap

Chapter 7 Health Management and Safety Assessment

7.1 Health Management Fundamentals

7.2 Unified Fault Modeling

7.3 Safety Reinforcement and Fault Data

7.4 Fault Hazards and Safety Assessment

7.5 Diagnosis and Prognostics

7.6 Safety Decisions and Improvement

7.7 Fault Injection Interfaces

7.8 Experiments and Validation

7.9 Exercises and Classroom Discussion

4 Units in This Section

7.1.1 Background and Significance

Textbook: Sec. 7.1.1

7.1.2 Health and Safety: Differences and Connections

Textbook: Sec. 7.1.2, 7.1.2.3

7.1.3 Health Management of Unmanned Systems

Textbook: Sec. 7.1.2.1

7.1.4 Algorithm Development Challenges and Validation

Textbook: Sec. 7.1.1

7.1.3 Health Management of Unmanned Systems: OSA-CBM

Six Functions Turn Measurements into Maintenance and Operating Advice

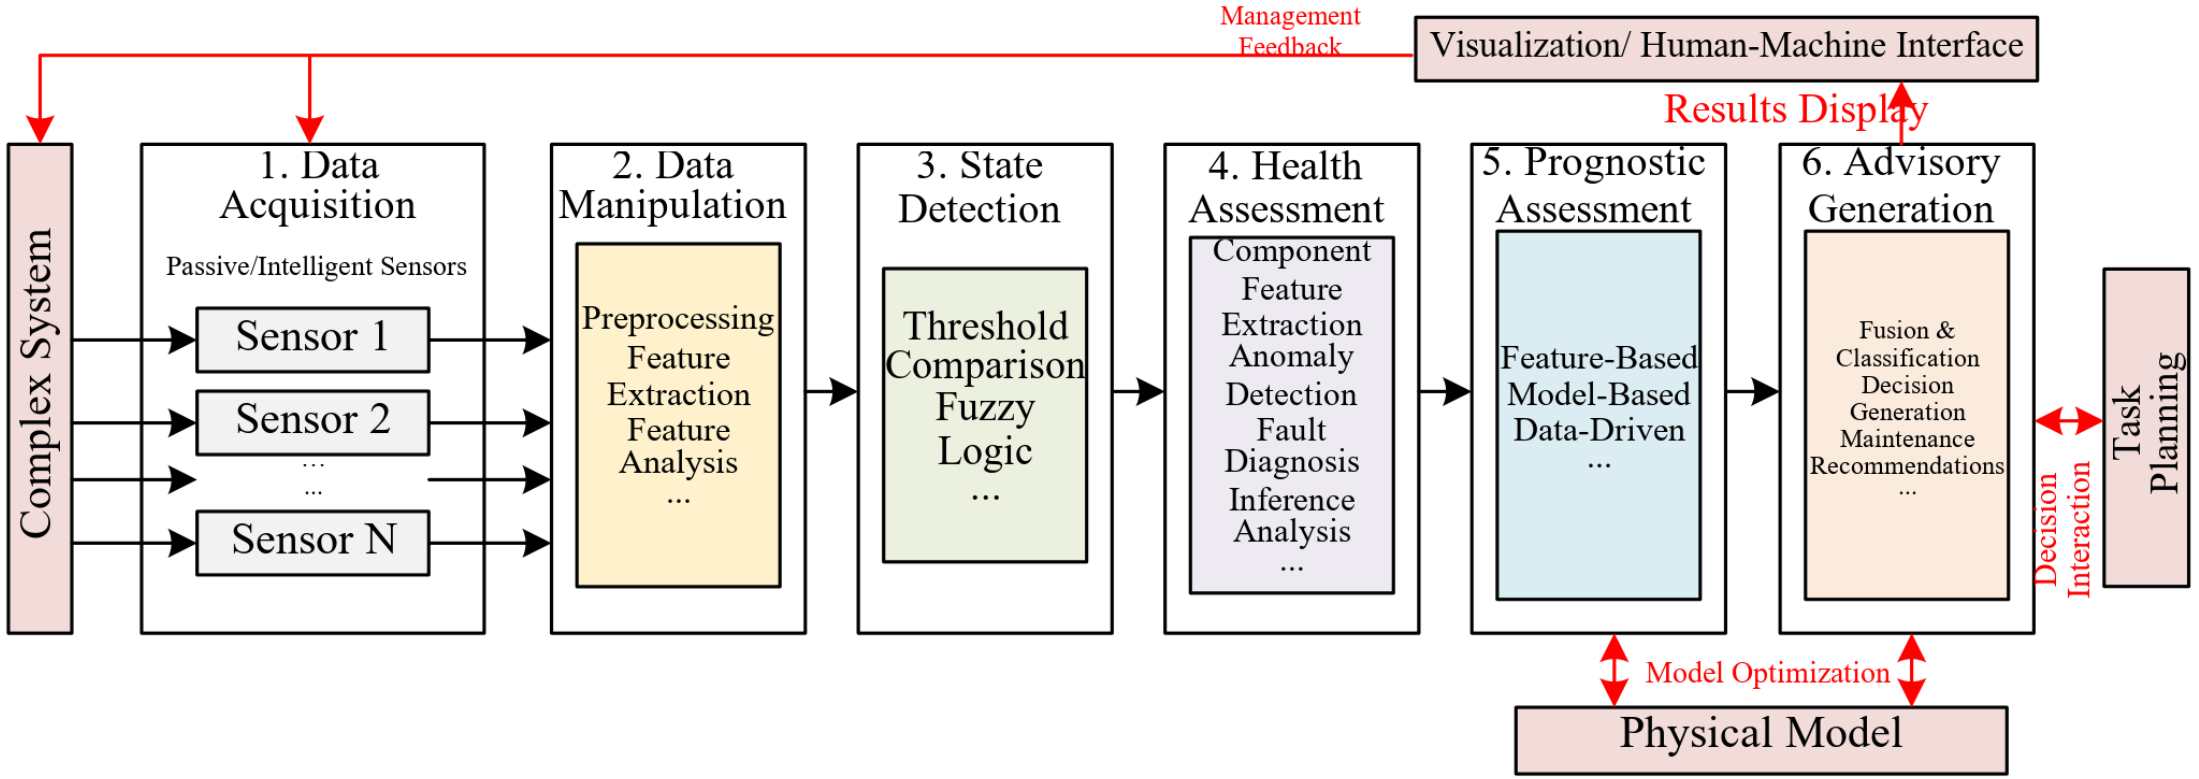


Fig. 7.1 Open System Architecture for Condition-Based Maintenance (OSA-CBM) for complex systems: six core PHM functions

7.1.3 Health Management of Unmanned Systems: Observation Chain

Acquisition, Processing, and Monitoring: Check the Observations

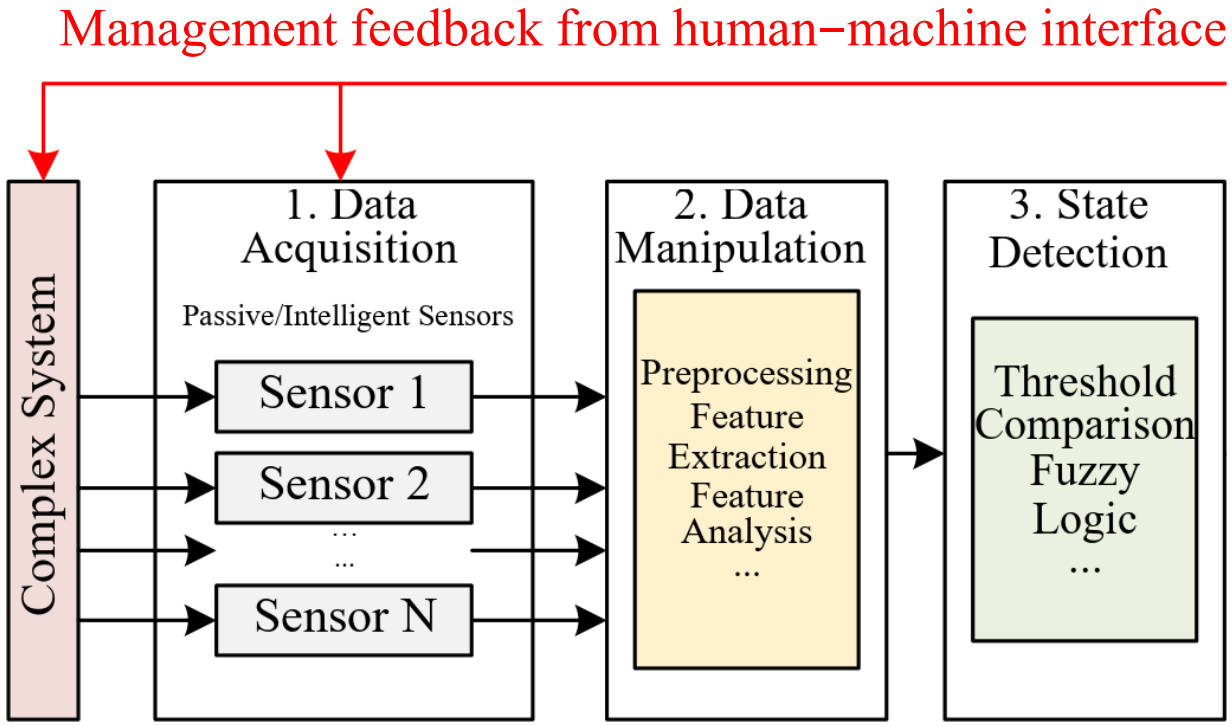


Fig. 7.1 (left): acquisition, processing, monitoring

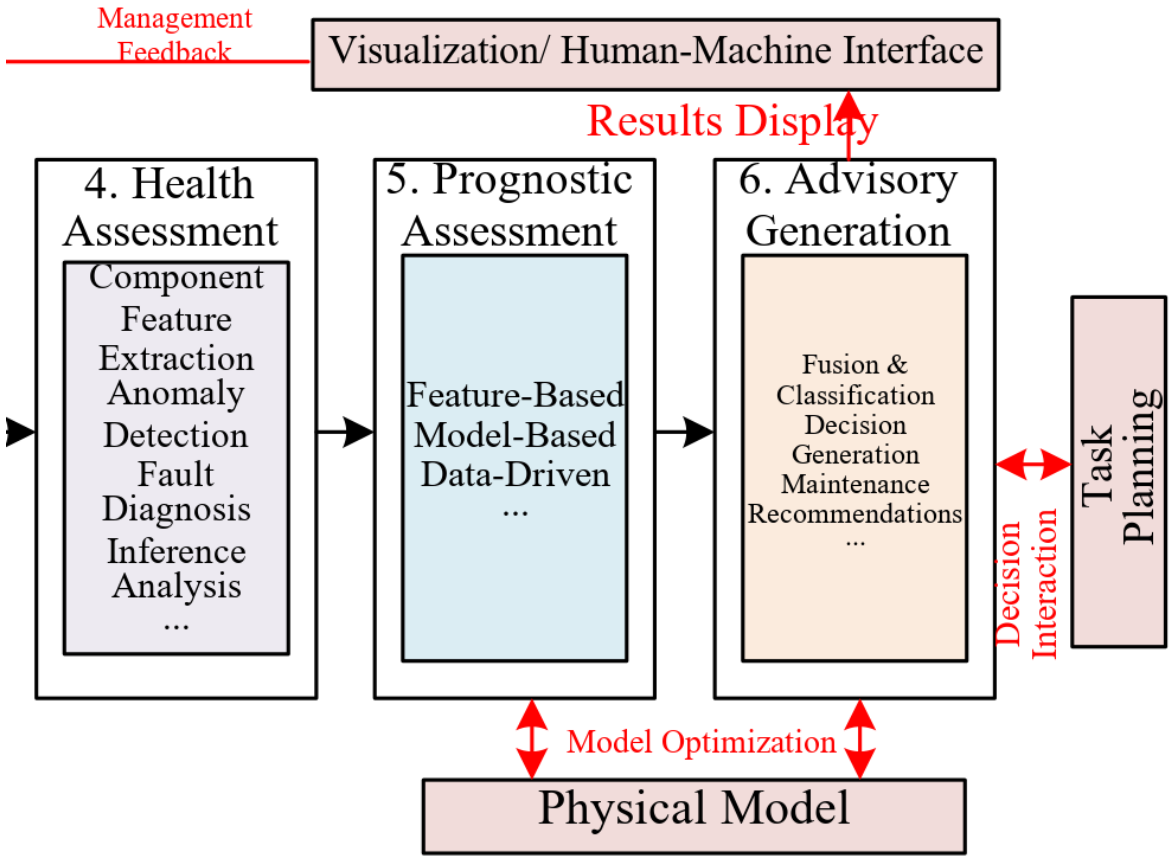
1. Data acquisition: current, speed, vibration, and control commands; retain timestamps and units.

2. Data manipulation: align time, segment by condition, extract features, preserve bias.

3. State detection: compare with a reference for the same conditions; detect anomalies and record events.

7.1.3 Health Management of Unmanned Systems: Assessment and Action

Diagnosis Explains the Present; Prognostics Informs Action



4. Health assessment: diagnose fault location, type, severity, and confidence.

5. Prognostic assessment: estimate degradation and life intervals for specified future conditions.

6. Advisory generation: recommend actions within mission constraints; reassess after execution.

Fig. 7.1 (right): health assessment through advisory generation

7.1.3 Health Management of Unmanned Systems: Data to Advice

One Motor Record Produces Different Information at Each Stage

Stage	Example input or output	Downstream context
Measurement and processing	Current, speed, vibration; windows under the same conditions	Time, units, missing-data rate, and processing method
Detection and diagnosis	Deviation events; evidence of reduced propulsion efficiency	Fault type, severity, and uncertainty
Prognostics and advice	Remaining capability; shorter-mission or maintenance advice	Future conditions; limits; validity period

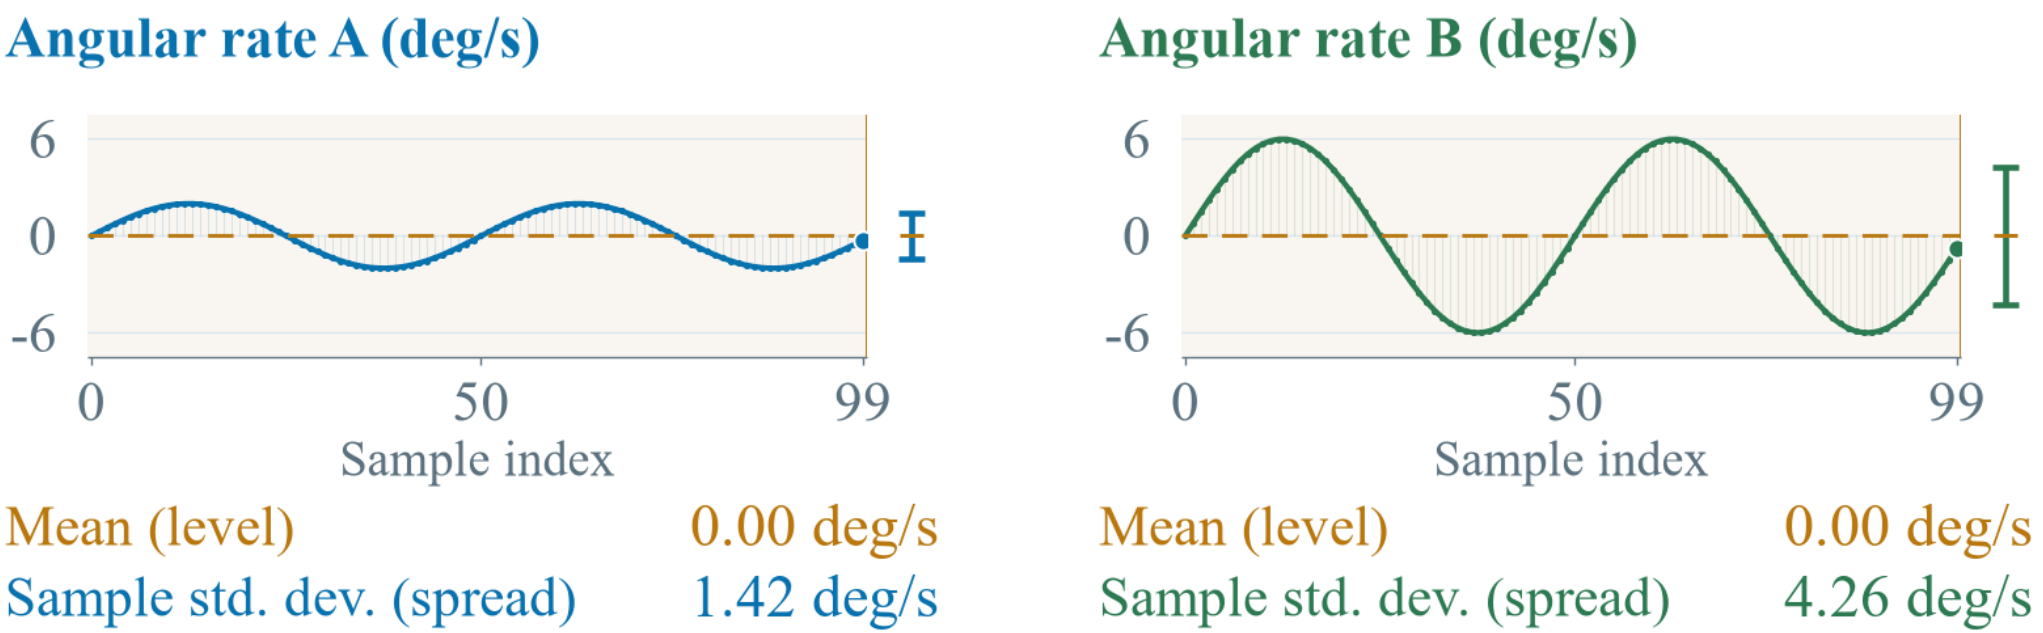
Health management provides conditional evidence for action, beyond an anomaly label.

7.1.3 Health Management of Unmanned Systems: Windows and Features

Two Signals Can Share a Mean but Differ Greatly

Same mean, different spread

100 Hz Window 100/100



Mean (level)	0.00 deg/s	Mean (level)	0.00 deg/s
Sample std. dev. (spread)	1.42 deg/s	Sample std. dev. (spread)	4.26 deg/s

Figure: in one window, the mean describes level; standard deviation describes fluctuation.

7.1.3 Health Management of Unmanned Systems: Residual Example

Residuals Compare Measurements with a Healthy Reference

Same conditions: healthy reference $\hat{y} = 6000\,\mathrm{r/min}$; measurement $y = 5400\,\mathrm{r/min}$.

Residual $$r = y - \hat{y} = 5400 - 6000 = -600\,\mathrm{r/min}$$

Relative residual $$r_{\mathrm{rel}} = \frac{y - \hat{y}}{\hat{y}} = -0.10$$

Interpretation: speed is 10% below the healthy reference; consider voltage, load, and other measurements.

Validation: build the reference from healthy data; use independent healthy and fault data to evaluate thresholds, false alarms, and missed detections.

7.1.3 Health Management of Unmanned Systems: Time and Decision Delay

Windows Reveal Detail but Also Require Waiting

Setting	Benefit	Cost
100 Hz sampling; 100-point window	Statistics over about 1 s	Wait for samples; sharp changes may blur
Slide every 0.1 s	More frequent assessment updates	Overlapping windows; dependent samples
Confirm persistent anomalies	Suppress false alarms from isolated noise	Longer confirmation can delay the response

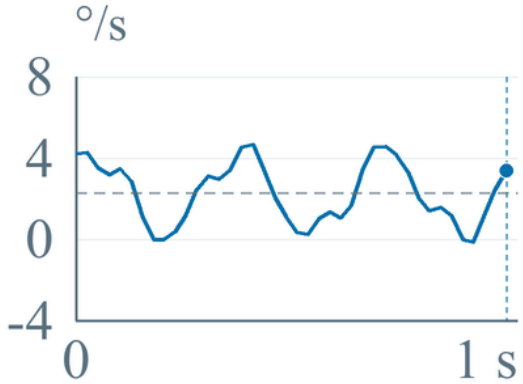
Beyond accuracy: can a diagnosis arrive in time to change a fault's consequences?

7.1.3 Health Management of Unmanned Systems: Classification Evidence

Extract Features, then Compare with Known States

Combine features; check borderline cases

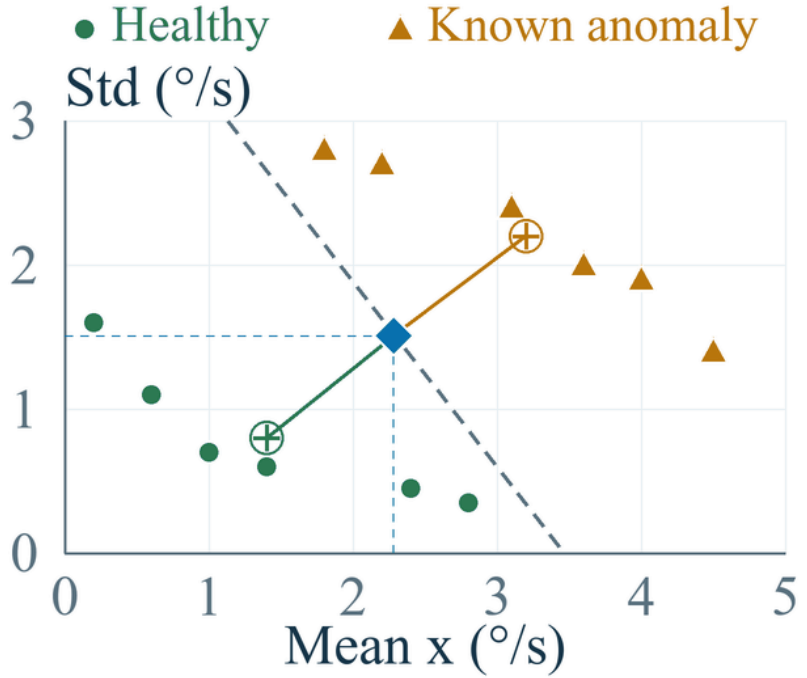
Synthetic angular rate



Mean x = 2.28 °/s

Std = 1.51 °/s

Window to feature point



New sample C

Distance to centers

Unit: °/s

Healthy 1.13

Anomaly 1.15

Similar distances

Seek more evidence

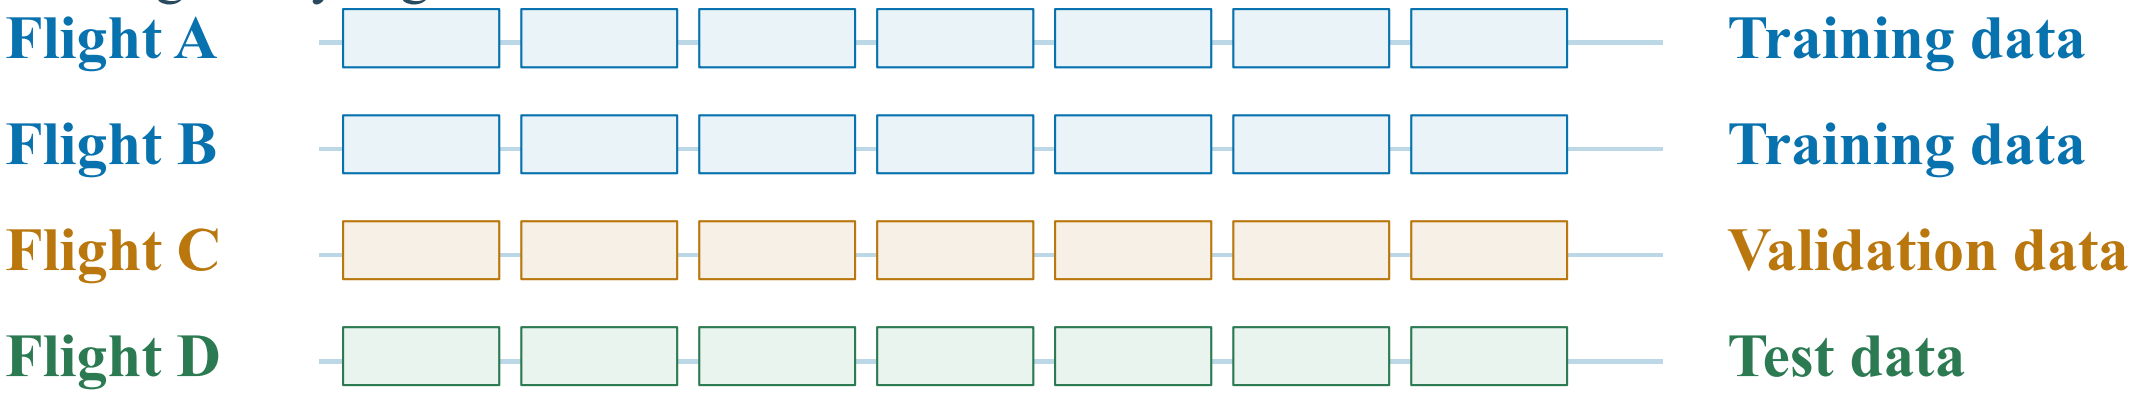
⊕ Class center

Figure: compare with known labels; new samples enter the classification space using the same features.

7.1.3 Health Management of Unmanned Systems: Independent Validation

Is Testing Independent if the Model Saw Neighboring Data?

Split whole flights first, then form windows within each group; neighboring windows from one flight stay together.



Offline Training: Features and Labels

Labels: verified states/fault records, time-aligned for learning and evaluation.

Online: available observed features.

Independent Test: Use Held-Out Flights

Fit scaling and model on training data.

Tune on validation; test only at the end.

Check false alarms, misses, and delay.

Chapter Outline

Health Management and Safety Assessment: Learning Roadmap

Chapter 7 Health Management and Safety Assessment

7.1 Health Management Fundamentals

7.2 Unified Fault Modeling

7.3 Safety Reinforcement and Fault Data

7.4 Fault Hazards and Safety Assessment

7.5 Diagnosis and Prognostics

7.6 Safety Decisions and Improvement

7.7 Fault Injection Interfaces

7.8 Experiments and Validation

7.9 Exercises and Classroom Discussion

4 Units in This Section

7.1.1 Background and Significance

Textbook: Sec. 7.1.1

7.1.2 Health and Safety: Differences and Connections

Textbook: Sec. 7.1.2, 7.1.2.3

7.1.3 Health Management of Unmanned Systems

Textbook: Sec. 7.1.2.1

7.1.4 Algorithm Development Challenges and Validation

Textbook: Sec. 7.1.1

7.1.4 Algorithm Development Challenges and Validation: Challenges

Five Challenges: Combinations, Data, Timing, Transfer, and Metrics

Many combinations: fault location, severity, time, conditions, and environment interact; one hover case is insufficient.

Scarce real fault data: severe faults are costly; simulation coverage and physical evidence must complement each other.

Limited real-time resources: diagnosis, data delay, and decision switching all consume the response window.

Conditional transfer: recheck models and thresholds after changing the vehicle, sensors, or mission.

Inconsistent metrics and tools: no unified assessment metrics; tools are difficult to reuse across different platforms.

7.1.4 Algorithm Development Challenges and Validation: Operating Levels

Control, Diagnosis, and Prognostics Use Different Time Scales

Motion Control	Measurement → estimation → control allocation → actuation Typically milliseconds (hundreds of Hz) for timely motion correction
Fault Diagnosis	Data window → anomaly detection → type / severity Typical windows: 0.1–1 s; balance evidence accumulation and detection delay
Prognostics and Decisions	Future capability + task limits → degraded mode / return / maintenance Typically minutes to flights; mission decisions respond to new risks

Emergency protection meets its own risk deadline without waiting for life prediction; levels share timestamped information.

7.1.4 Algorithm Development Challenges and Validation: Evidence Levels

Simulation Extends Coverage; Physical Tests Check Model Omissions

Stage	Main evidence	Still needed
Offline records	Data quality, detection rate, false alarms, and delay	Consequences after executing actions
Software-in-the-loop (SITL; SIL in textbook)	Closed-loop response across many fault combinations	Model error and real hardware constraints
Hardware-in-the-loop (HITL; HIL in textbook)	Computing load, links, and timing effects	Physical propulsion and sensing differences
Controlled physical tests	Actual response under specified conditions	Conditions outside test coverage

7.1.4 Algorithm Development Challenges and Validation: Platform Support

Five Ways RflySim Supports Health Management and Safety Assessment

Batch extreme-condition tests: extreme weather, satellite-navigation interference, communication loss, and sudden battery depletion.

SITL / HITL closed-loop assessment: SITL checks functionality; HITL tests delays and hardware disturbances.

State monitoring and visualization: observe state variables, health indicators, and diagnostic outputs in simulation.

Data export and training: export through Python, Simulink, and Robot Operating System (ROS) for offline model training.

Simulation-to-real transfer: deploy algorithms to the real vehicle by automatic code generation, then check consistency.

Exercises and Discussion

Exercise: Small Error, Hot Motor—What Evidence Comes Next?

Observed: height error stays within 0.1 m; one motor heats up; the control command is 20% above the healthy reference.

Question 1: which statements are observations, and which are hypotheses about causes?

Question 2: what further data distinguish propulsion degradation, a load change, and wind disturbance?

Question 3: why is a small height error insufficient reason to continue the original inspection mission?

Reference Points

Gather Evidence, then Constrain the Mission by Remaining Capability

Facts: temperature and command changes. Hypotheses: motor fault, gusts, or changed load.

More data: speed, current, voltage, load, and wind; use a healthy reference for the same conditions.

Distinguish causes: one command rising or all four? Relate heating to wind.

- **Thrust effectiveness loss:** higher command for equal thrust; interpret current and heating by mechanism.
- **Added load:** all four commands rise; motor heating is similar.
- **Wind:** commands vary with heading/wind, with tilt; recover when wind stops.

Decision: check control margin, return energy, and landing sites; shorten the mission or activate protection if needed.

Pitfall: small errors can reflect compensation; normal tracking does not establish health.

Section Summary

From Detecting an Anomaly to Knowing the Remaining Capability

Concepts: anomaly, fault, and failure describe different levels; degradation may be compensated or exhaust control capability.

Assessment: health describes condition and capability; safety also depends on the mission, environment, exposed people and property, and safeguards.

Method: acquisition, processing, and diagnosis lead to prognostics and advice; check results using offline records, SITL, HITL, and controlled physical tests.

Common errors: treating symptoms as causes, current health as mission safety, or one simulation level as complete evidence.

Next: model possible faults to make injection, comparison, and improvement repeatable.