



AF-TR-9600-001

Quantitative Methods and Worked Calculations

Space / Aviation / Defense

FIRST EDITION

AF 9600 SUPPORTING DOCUMENT

Document identity and use

FIRST EDITION MANUSCRIPT - Prepared for Aerospace Federation review

Field	Controlled identity
Document	AF-TR-9600-001 - Quantitative Methods and Worked Calculations
Parent standard	AF 9600 - Artificial Intelligence, Autonomous Systems & Interoperability
Edition / language	FIRST EDITION / English master text
Family	Companions
Primary source anchors	6, 9, 15, 19, 22-25, 34; informative Annex O
Suggested local owner	Systems / statistics / assurance specialist
Authority	Publisher approval and effective date are not assigned. Local use requires authorized adoption and completed parameters.

Use with the main standard

This document implements the identified AF 9600 clauses and annexes. Read the complete applicable requirements in the main standard. The source anchors are navigation references; a template does not replace technical judgement, evidence, permission or eligible approval.

AF-ML-9600-001 lists the complete supporting set. AF-CM-9600-001 maps the main clauses and annexes to working documents. AF-TP-9600-001 establishes the common local parameter baseline; add the relevant discipline profiles.

Record control

Identify the object, configuration, record owner and dates. Reference attachments by stable identifier and version. Keep superseded evidence and decisions traceable. Use the approved local retention, access and authentication arrangements. No technical limit, numerical limit, interval or acceptance value is assigned by a blank template.

Implementation model: Translate the stated responsibilities and controls into the actual organization. Informative examples remain explanatory; use approved programme-specific values for operational decisions.

1 Use of quantitative examples

All values and cases in this companion are synthetic. They explain assumptions, calculations and limitations; they are not system limits, product evidence or mandatory default parameters. Select and approve the actual model and values before relying on a calculation.

2 Worked calculations

Status: Informative. All values are synthetic and are not mandatory thresholds, actual product results or approved operating parameters. Each example states assumptions and the limit of its conclusion.

0.1 Aggregate accuracy and critical populations

An independent evaluation contains 9,900 ordinary cases, of which 9,801 are correct, and 100 rare consequential cases, of which 60 are correct. Assume the labels and the specified sampling design are valid.

Ordinary accuracy = $9,801 / 9,900 = 99.00\%$

Rare-case accuracy = $60 / 100 = 60.00\%$

Aggregate accuracy = $(9,801 + 60) / 10,000 = 98.61\%$

Rare-case error fraction = $40 / 100 = 40.00\%$

The aggregate conceals substantial failure in the rare population. It cannot compensate for an unmet rare-case criterion. Whether the sample proportions represent operation must be assessed separately; applying operational weights changes an aggregate, not the observed subgroup result.

0.2 Independent units and zero observed failures

An evaluation includes 200 independently sampled missions with 50 highly related frames per mission. No mission contains the defined failure. Assume missions are independent Bernoulli trials with a constant mission-failure probability and reliable detection. The exact one-sided 95% upper confidence bound for that probability is:

Observed frames = $200 \times 50 = 10,000$

Independent mission trials = 200

Upper bound = $1 - 0.05^{(1 / 200)} = 0.0148670$

Upper bound as percentage = approximately 1.4867%

Using 10,000 as the independent trial count would be unjustified. The result concerns the defined event per mission in the sampled population. It is not a universal failure rate, a probability that software is safe, or evidence that unrepresented mission conditions are acceptable.

0.3 Selective prediction and fallback demand

A system receives 10,000 cases. It accepts 8,000 for automated prediction and abstains on 2,000. Of accepted cases, 7,600 are correct. Assume all cases and outcomes are retained.

Automated coverage = $8,000 / 10,000 = 80\%$

Accuracy conditional on acceptance = $7,600 / 8,000 = 95\%$

Correct automated answers / all cases = $7,600 / 10,000 = 76\%$

Abstention fraction = $2,000 / 10,000 = 20\%$

If the input rate is 500 cases per hour and proportions remain applicable, expected fallback demand is $500 \times 0.20 = 100$ cases per hour. This is an average demand, not a peak-capacity guarantee. Human

review latency, burst arrivals and correlated difficult cases require separate assessment.

0.4 Calibration error within declared bins

Two equal-size bins contain 100 predictions each. Their mean predicted probabilities are 0.90 and 0.60; observed event fractions are 0.80 and 0.70. For this illustrative two-bin expected calibration error:

$$\text{Bin 1 absolute gap} = \text{abs}(0.90 - 0.80) = 0.10$$

$$\text{Bin 2 absolute gap} = \text{abs}(0.60 - 0.70) = 0.10$$

$$\text{ECE} = (100 / 200) \times 0.10 + (100 / 200) \times 0.10 = 0.10$$

The reported error is 10 percentage points for these bins and data. Bin selection, small samples and subgroup variation can materially change interpretation. Even perfect calibration of the defined event would not establish correct action authority, complete system safety or valid operation outside the evaluated population.

0.5 Protective margin for an illustrative ground cart

A ground inspection cart travels at 2.0 m/s toward a defined limit. Assume conservative sequential bounds of 0.08 s sensing, 0.05 s inference, 0.02 s monitoring, 0.05 s switching and 0.10 s actuator delay. After that delay, assume constant deceleration of at least 1.5 m/s squared. Add a 0.20 m position-error allowance and a 0.30 m design reserve.

$$\text{Response delay} = 0.08 + 0.05 + 0.02 + 0.05 + 0.10 = 0.30 \text{ s}$$

$$\text{Travel during delay} = 2.0 \times 0.30 = 0.60 \text{ m}$$

$$\text{Idealized braking distance} = 2.0^2 / (2 \times 1.5) = 1.333333 \text{ m}$$

$$\text{Required distance} = 0.60 + 1.333333 + 0.20 + 0.30 = 2.433333 \text{ m}$$

$$\text{Available distance} = 2.80 \text{ m}$$

$$\text{Remaining margin} = 2.80 - 2.433333 = 0.366667 \text{ m}$$

This model excludes slopes, jerk limits, changing friction, actuator faults and obstacles entering the path. Those effects must be included or bounded before operational use. The example is not a flight-control rule or a universal stopping distance. An average measured delay is not a conservative bound.

0.6 Human intervention feasibility

Assume a task has a justified 3.0 s intervention window. The proposed response chain uses bounds of 0.4 s detection, 0.3 s notification, 2.0 s recognition, 1.2 s decision and 0.6 s actuation.

$$\text{Complete intervention} = 0.4 + 0.3 + 2.0 + 1.2 + 0.6 = 4.5 \text{ s}$$

$$\text{Window shortfall} = 4.5 - 3.0 = 1.5 \text{ s}$$

The proposed human intervention does not fit the stated window. A visible alert does not close this gap. The design needs a demonstrably faster protective path, a larger justified margin, or restricted operation. Human-time assumptions must be validated for representative users, workload and abnormal conditions; isolated best-case responses cannot establish these bounds.

0.7 Data age with clock uncertainty

A source records an event at 100.000 s with an absolute clock-error bound of 0.010 s. A consumer evaluates the data at 100.120 s with an error bound of 0.015 s. Assume a shared time convention and that both bounds cover the relevant capture instants.

$$\text{Source event interval} = [99.990, 100.010] \text{ s}$$

$$\text{Consumer use interval} = [100.105, 100.135] \text{ s}$$

$$\text{Possible data age} = [100.105 - 100.010, 100.135 - 99.990]$$

Possible data age = [0.095, 0.145] s

For a locally approved maximum age of 0.120 s, the nominal 0.120 s difference does not demonstrate compliance: the upper bound is 0.145 s. The consumer must apply its approved uncertainty-aware rule and protective behavior. These are interval bounds, not standard deviations to combine by root-sum-square.

0.8 Units and coordinate transformation

A two-dimensional source supplies a point (100, 200) in centimetres. The approved transform first converts to metres, then rotates the source frame by +90 degrees counterclockwise into the destination frame, then adds translation (10, 20) metres. Assume the stated axis order and orientation are correct.

Source in metres = (1, 2)

Rotation matrix = $\begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$

Rotated point = (-2, 1) m

Destination point = (-2 + 10, 1 + 20) = (8, 21) m

Translating before rotating or interpreting centimetres as metres produces a different result even though the payload can pass a schema check. Actual systems must also account for transformation uncertainty, time dependence, calibration and the validity of the reference frames.

0.9 Collective communication capacity

Forty participants each send 20 messages per second with a 300-byte payload. Assume a measured overhead factor of 1.40 for this traffic pattern and a 5 Mbit/s link with an approved usable allocation of 60%. Use decimal bits and bytes.

Payload rate = $40 \times 20 \times 300 = 240,000$ bytes/s

Payload bit rate = $240,000 \times 8 = 1,920,000$ bit/s

Estimated carried rate = $1,920,000 \times 1.40 = 2,688,000$ bit/s

Usable allocation = $5,000,000 \times 0.60 = 3,000,000$ bit/s

Remaining allocation = $312,000$ bit/s = 0.312 Mbit/s

Average capacity headroom is not a deadline proof. Bursts, discovery, retries, management traffic, priority inversion and queueing need separate evidence. The overhead factor is an example-specific assumption and shall not be reused without measurement or a justified model.

0.10 Compatibility matrix coverage

A programme requires three producer versions, four consumer versions and two operating modes, giving 24 directed version-mode combinations. Eighteen combinations have valid tests: 16 pass and two fail. Six are untested.

Required combinations = $3 \times 4 \times 2 = 24$

Tested coverage = $18 / 24 = 75\%$

Passing combinations / required combinations = $16 / 24 = 66.6667\%$

Passing combinations / tested combinations = $16 / 18 = 88.8889\%$

The last percentage must not be presented as complete interoperability. Required failed or untested relationships remain blocked unless an eligible, explicitly restricted scope excludes their use. Directionality, feature options and additional partners may enlarge the matrix beyond this simple product.

O.11 Zero events over exposure time

An observation programme records zero events of a defined type during 10,000 hours. Assume a homogeneous Poisson process, complete event detection and representative exposure. The one-sided 95% upper bound on the event intensity is:

$$\text{Upper intensity} = -\ln(0.05) / 10,000$$

$$\text{Upper intensity} = 0.000299573 \text{ events/hour}$$

This is an intensity bound under the specified model, not a universal safety claim. Dependence, changing conditions, incomplete reporting or an incorrectly defined event invalidate the simple interpretation. Hours of simulation cannot be substituted for operational hours without a justified relationship.

O.12 Correlated estimates and fusion uncertainty

Two unbiased estimates of the same scalar state have standard deviations 2 and 3 units and correlation 0.5. Form their equal-weight average. Assume finite variances, correct covariance and no unmodeled common bias.

$$\text{Variance} = 0.5^2 \times 2^2 + 0.5^2 \times 3^2$$

$$+ 2 \times 0.5 \times 0.5 \times 0.5 \times 2 \times 3 = 4.75$$

$$\text{Standard deviation} = \sqrt{4.75} = 2.179449 \text{ units}$$

$$\text{If independence were wrongly assumed: } \sqrt{3.25} = 1.802776 \text{ units}$$

Ignoring correlation understates uncertainty in this example. This calculation does not prove optimal fusion or remove bias. Repeated information from one original sensor through multiple interfaces may be even more strongly dependent and requires explicit provenance analysis.