On the Applicability of Safety Nets: A Safety-By-Design Solution for Certifying Neural Networks
ICAS PAPER 2026_0210



Johann Christensen, Thomas Stefani, Elena Hoemann, Frank Köster, Sven Hallerbach
Institute for AI Safety and Security, German Aerospace Center (DLR)



Motivation

$10 bn

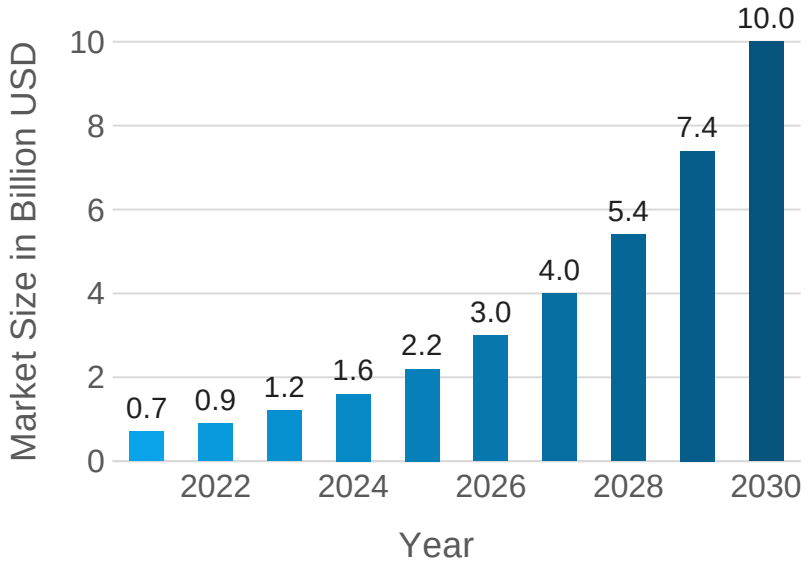
AI-in-aviation market by 2030

700 000

pilots short through 2043

Level 1–2 AI

reaching the cockpit this decade



Take-away More AI in the cockpit is coming. The open question is how to certify it.



Precedence Research, “Artificial Intelligence in Aviation Market Size, Share, and Trends 2024 to 2034”. Available: https://www.precedenceresearch.com/artificial-intelligence-in-aviation-market

Boeing, “2024 Pilot and Technician Outlook”. Available: https://boeing.com/pto

Three Take-Aways

1 Correct by construction

Network + lookup table = 100 % agreement with ground truth over the whole ODD.

2 Architecture sets the price, not the safety

Every architecture is equally certifiable—it only influences the size.

3 Small enough to fly, open enough to check

4.76 MiB and 222.83 MiB instead of ≥ 4 GiB, fully reproducible.



ACAS X: Two Walls

MEMORY WALL

≥ 4 GiB

to embed the ACAS X logic tables

512 MB

RAM on current FMCs

CERTIFICATION WALL

EASA: Safety-by-Design

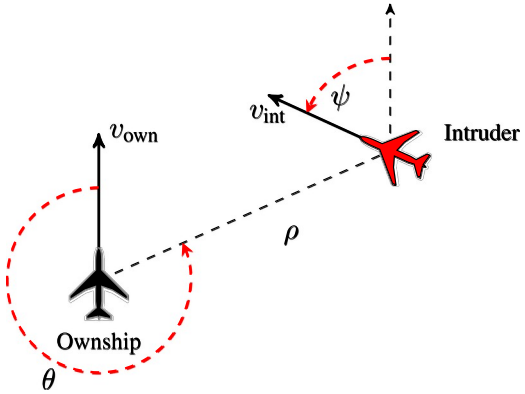
Mandatory for AI-based systems in safety-critical applications

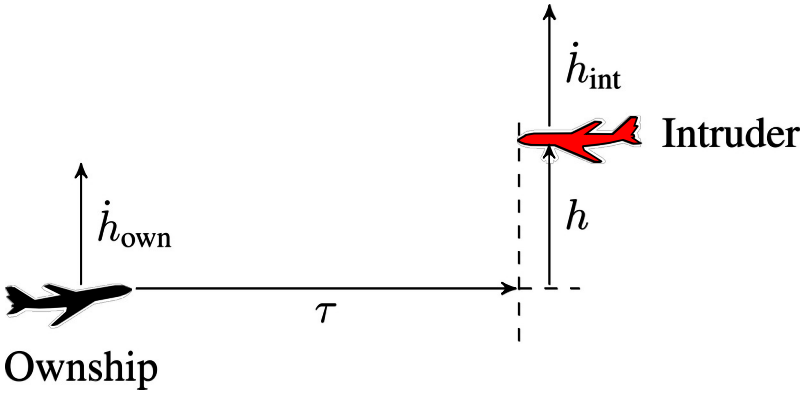
No finite test set can certify a complete continuous input space

Take-away Compression alone is not enough—it has to be provably correct.



Use Case: HCAS and VCAS





HCAS Horizontal Advisories

7 inputs · 5 advisories · 40 subsystems
53 792 states per subsystem
2 151 680 states in total

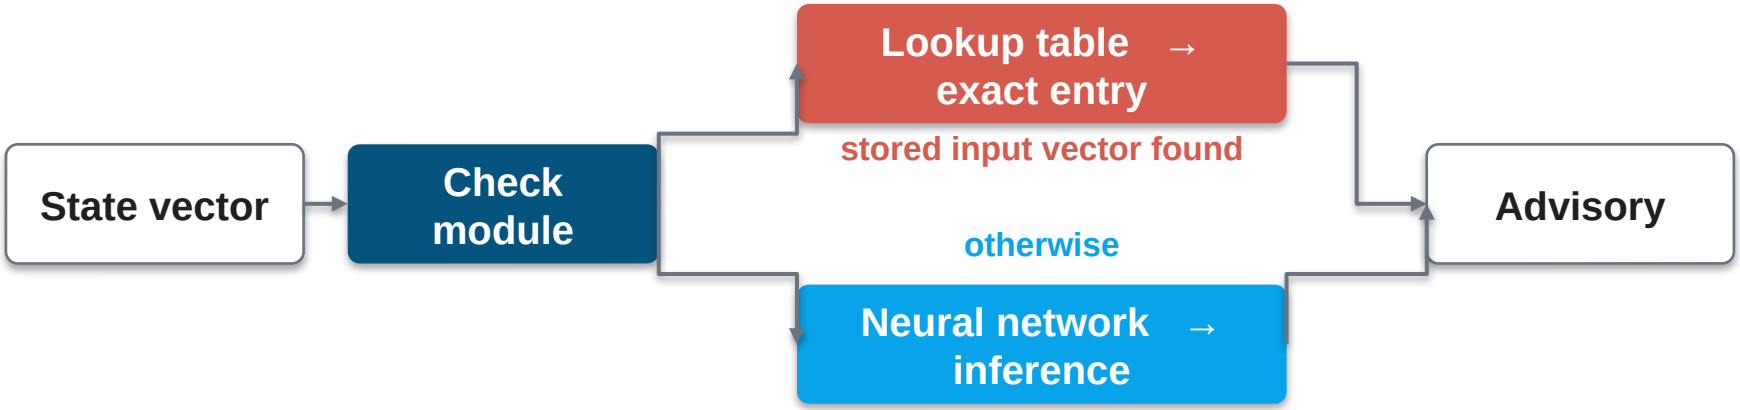
VCAS Vertical Advisories

5 inputs · 9 advisories · 9 subsystems
4 053 465 states per subsystem
36 481 185 states in total



K. D. Julian and M. J. Kochenderfer, “Guaranteeing Safety for Neural Network-Based Aircraft Collision Avoidance Systems,” in *2019 IEEE/AIAA 38th Digital Avionics Systems Conference (DASC)*, 2019.

The Safety Net





Take-away Better network → smaller table. Worse network → bigger table. Always 100 % correct.



Building and Verifying the Safety Net

HYPERPARAMETER SWEEP

Activation: ReLU · LeakyReLU · GELU

Hidden layers: 2 · 3 · 5 · 7

Nodes per layer: 25 · 50 · 100 · 150 · 200

Encoding: one-hot vs. target

5 runs per configuration
≤ 10 000 epochs · patience 1 000

NO HELD-OUT SPLIT

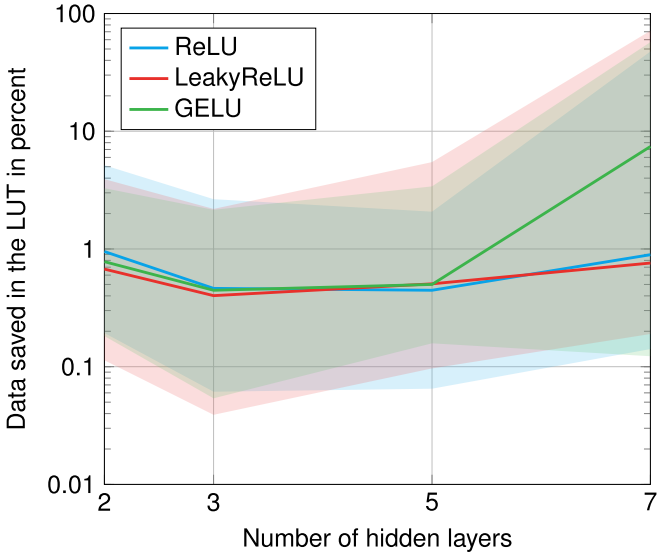
Finite, fully known input space—compression, not generalization.

BUILDING IS VERIFYING

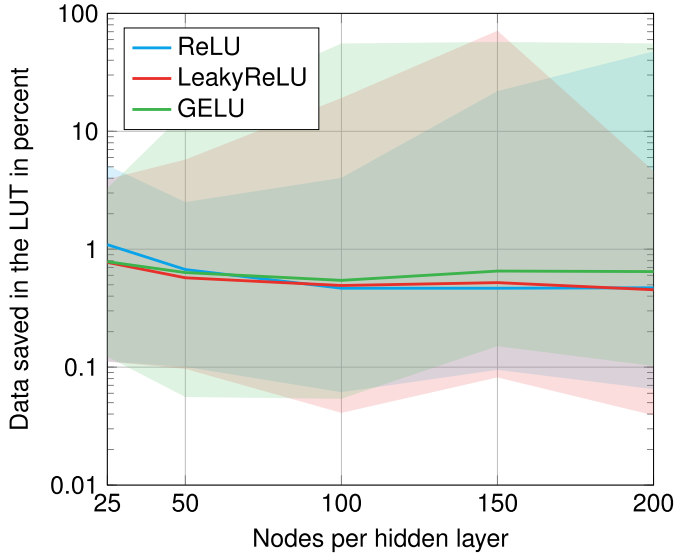
1. **Train the network**
2. **Sweep every input vector**
3. **Collect every deviation**
4. **Write the LUT**



Depth, Width, and Activation Function



Optimum at 3–5 hidden layers

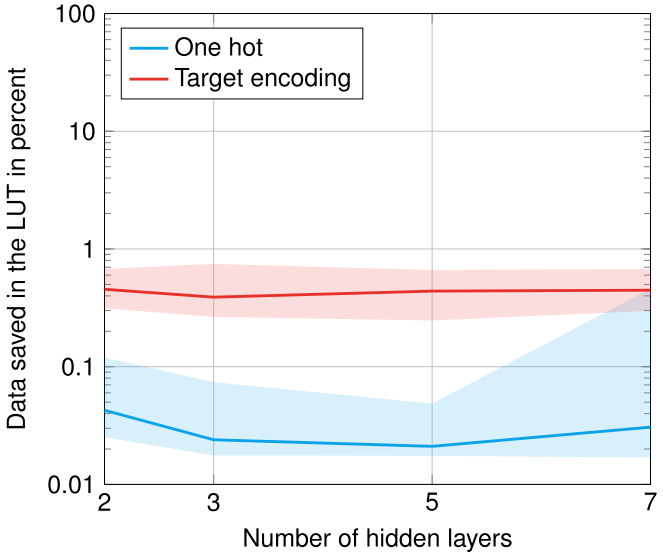


Saturation beyond 50–100 nodes

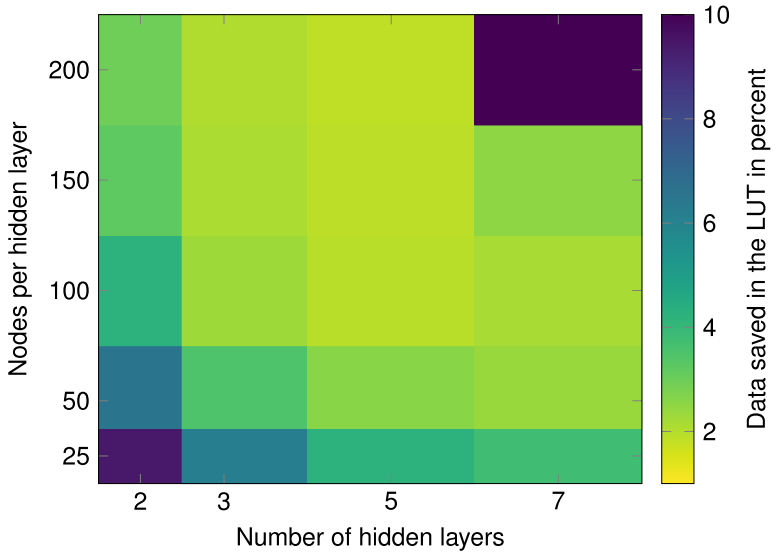
Take-away ReLU wins · depth and width buy variance, not accuracy.



Encoding and Robustness of the Choice



One-hot vs. target encoding (VCAS, ReLU)



ReLU across the whole grid (VCAS)

Take-away One-hot gains one order of magnitude · ReLU is forgiving.



The Resulting Safety Nets

System	Neural network	LUT	Combined	Network coverage
HCAS	3 MiB	2 MiB	**5 MiB**	99.7 %
VCAS	1 MiB	221 MiB	**223 MiB**	97.8 %

≈ 3 orders of magnitude
smaller than the 4 GiB table (HCAS)

20–46 µs
inference time against a 1 s update budget

HCAS: 3 hidden layers × 100 nodes · VCAS: 5 × 100 · ReLU, one-hot · LUT stored as a k-d tree

Take-away Both fit the 512 MB memory budget of current avionics hardware.



What This Buys for Certification

LM-04 **Quantifiable generalization bounds**

Statistical bounds are too loose; the sweep gives a deterministic zero gap.

LM-10 / IMP-11 **Requirements-based verification**

The lookup table is the specification; every input vector is checked against it.

LM-09, 12–14, 16 · IMP-09, 10, 12 **Stability, robustness, completeness**

By-products of the same sweep—no separate verification campaign.

Under the Proposed Issue 03: LM-04 → S-LM-06 + SU-LM-14 · LM-10 → SU-LM-11 · IMP-11 → SU-IMP-06/07



Limitations and Next Steps

LUT dominates VCAS

99.4 % of the combined VCAS size. Per-subsystem optimization should shrink it.

Point-wise, not interval

The guarantee holds on the discretized grid, not between grid points.

223 MiB → ≈ 20 MiB

Plain binary serialization instead of a k-d tree.

A reduced surrogate

Full envelopes are orders of magnitude larger.

Next Per-subsystem search · sin/cos angles · pruning and quantization



Conclusion

1 Correct by construction

100 % agreement with ground truth over the whole discretized ODD—a deterministic bound, not a statistical one.

2 Architecture sets the price, not the safety

Optimum: ReLU · 3–5 hidden layers · 50–100 nodes · one-hot encoding.

3 Small enough to fly, open enough to check

5 MiB and 223 MiB against ≥ 4 GiB—the first publicly reproducible Safety Nets for HCAS and VCAS.

Code DLR-KI/castrainer **Artifacts** 10.5281/zenodo.20666415 **Contact** johann.christensen@dlr.de

