The Anatomy of a Misprediction: Tracking Terrestrial and Orbital Radiation-Induced Errors in Systolic Arrays

Rafael Tonetto, Univ. Rennes; Pedro Foletto Pimenta, Univ. Trento; Abraham Chan, Karthik Pattabiraman, Univ. of British Columbia; Fernando Fernandes dos Santos, Marcello Traiola, Angeliki Kritikakou, Univ. Rennes; Paolo Rech, Univ. Trento

SC26 hpc unites | Chicago

Why should we care about a radiation-induced faults?

AI is moving into places where an incorrect output is not just a bad score.

AUTONOMOUS SYSTEMS

A silent corruption can change a perception or decision without producing a crash.

SPACE AI

Radiation exposure is dramatically higher in orbit; AI accelerators become part of the reliability problem.

THE HARD PART

The final wrong answer tells us almost nothing about where the particle struck or how the error propagated.

The central reliability question is not “can radiation flip a bit?”

It is: which hardware faults actually become critical AI failures?

The reliability gap: three layers, three different answers

BEAM EXPERIMENT

Real particles

✓ Real device
✓ Real physics

✗ Expensive
✗ Poor observability
✗ Hard to trace a failure back to a resource

RTL / GATE-LEVEL FI

Observable hardware

✓ Fault location
✓ Propagation
✓ Resource-level vulnerability

✗ Expensive at scale
✗ Proprietary accelerator details

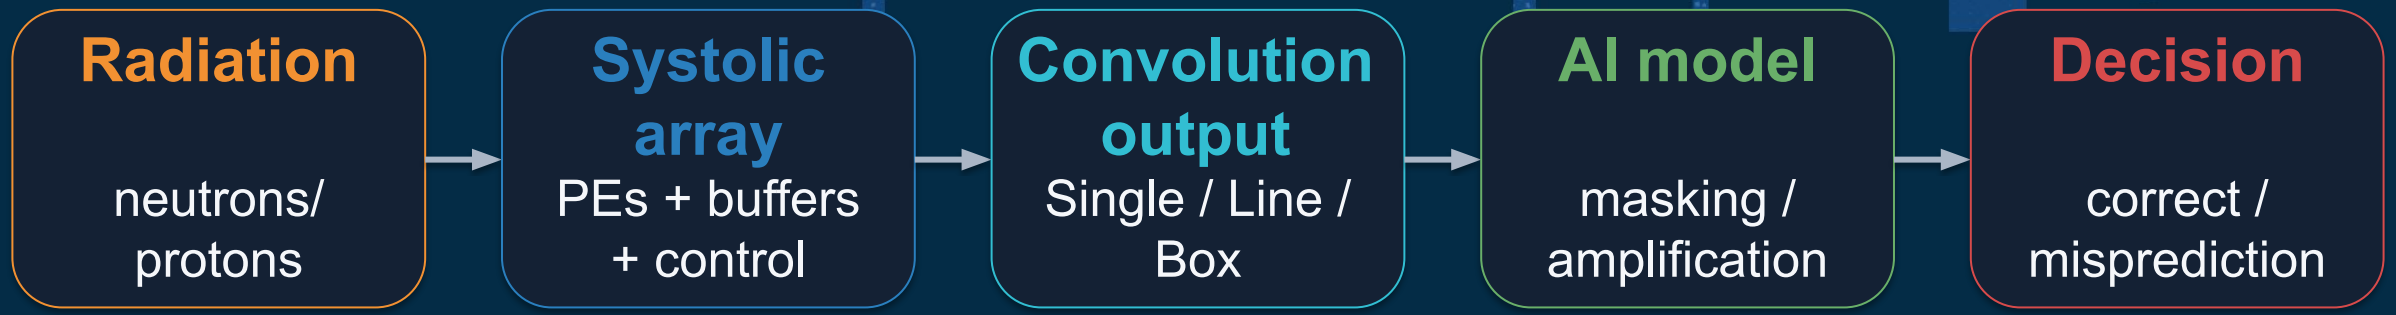
SOFTWARE FI

Fast and scalable

✓ Full models
✓ Layer-by-layer analysis
✓ Millions of injections
✗ What should be injected?
✗ Single-bit models can be unrealistic

The paper connects the three instead of choosing one.

One cross-layer question: particle→hardware→tensor→decision



Measure frequency physically → identify causes in RTL/GL → propagate realistic patterns in software

HW VULNERABILITY

HVF: probability that a hardware fault creates a given output pattern.

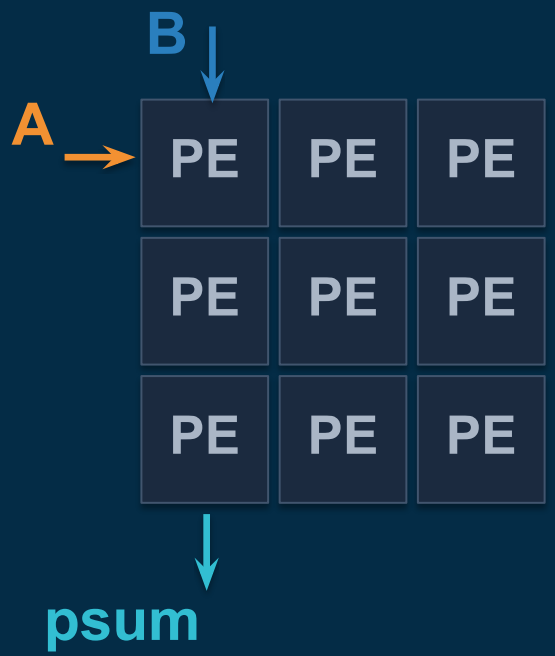
SW VULNERABILITY

PVF: probability that that pattern becomes a critical model error.

AVF = HVF × PVF

Why systolic arrays create a different fault model

Weight-stationary systolic array



LOCAL FAULT

MAC / partial sum → one output location

PERSISTENT FAULT

Weight / control → propagates through time and PEs

REUSED DATA

Input-buffer fault → same bad pixel enters multiple patches

RESULT

One particle can create a structured tensor corruption — not a single bad bit.

That geometry is the bridge between silicon physics and AI behavior.

Beam experiments on a Google's Coral TPU

7

Coral TPUs exposed

17 h

neutron beam time

24 h

proton beam time

>760k yr

Earth-equivalent exposure

>278 yr

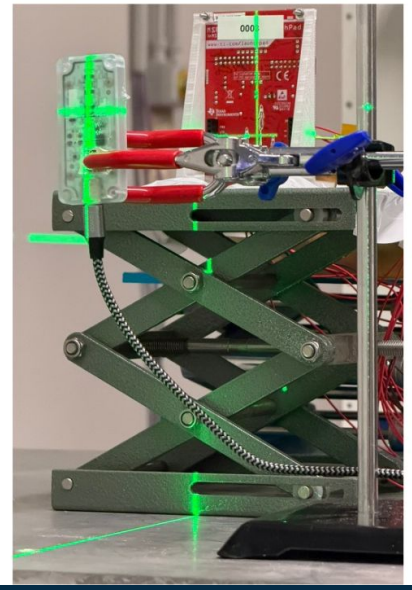
orbital-equivalent exposure

CHIPIR



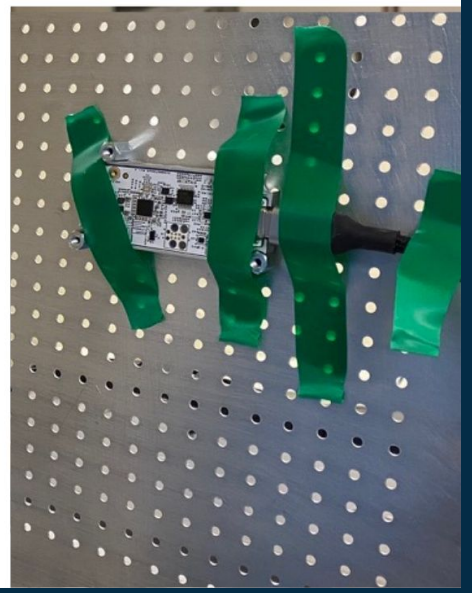
Neutrons

CNAO



Protons

PATREC



Protons

SC26

Beam experiments on a Google's Coral TPU

7

Coral TPUs exposed

17 h

neutron beam time

24 h

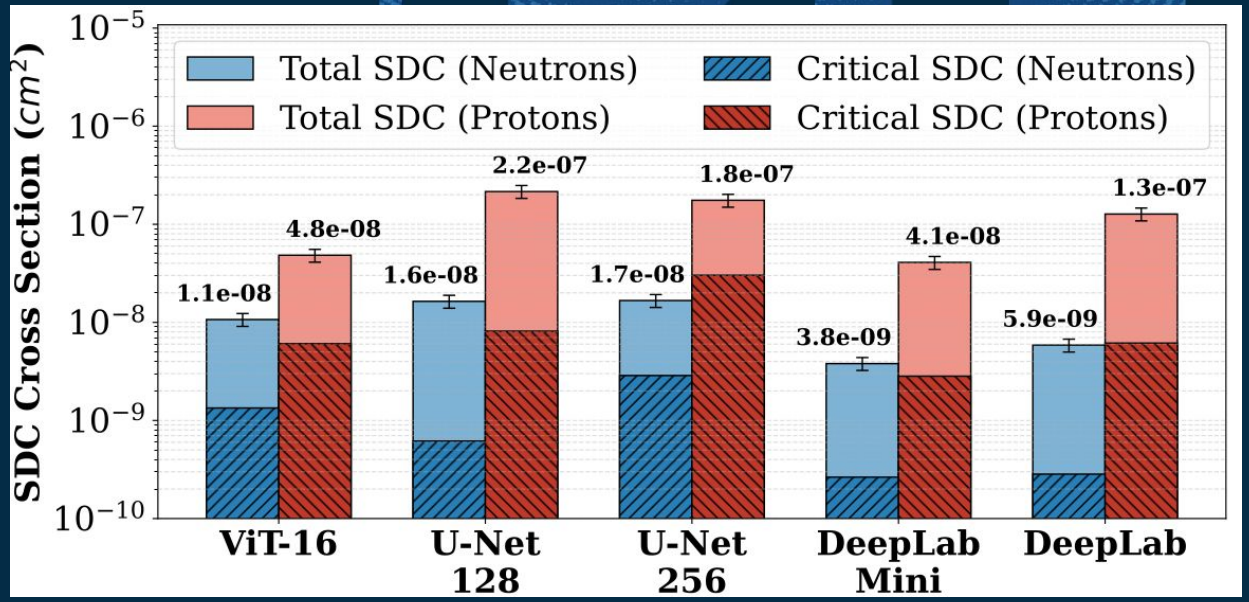
proton beam time

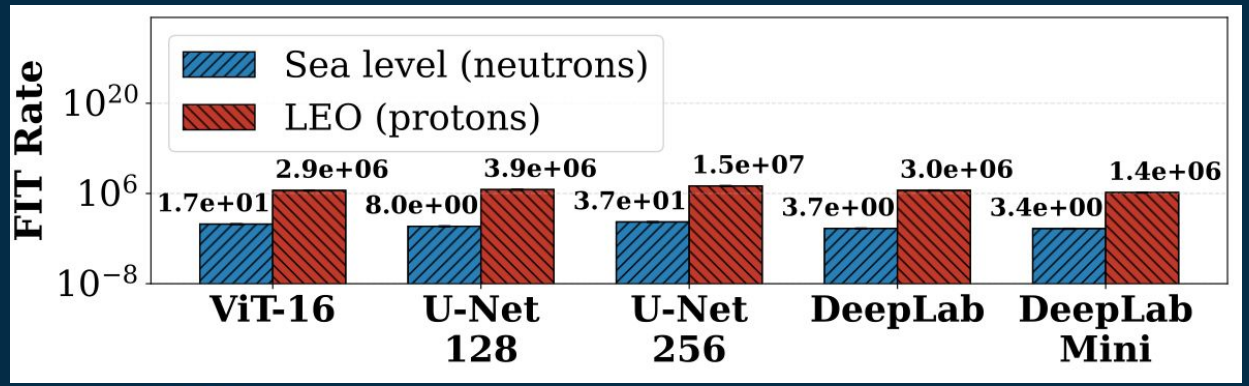
>760k yr

Earth-equivalent exposure

>278 yr

orbital-equivalent exposure





First result: orbit changes the reliability scale

200 MeV proton vs. neutron

11.3×

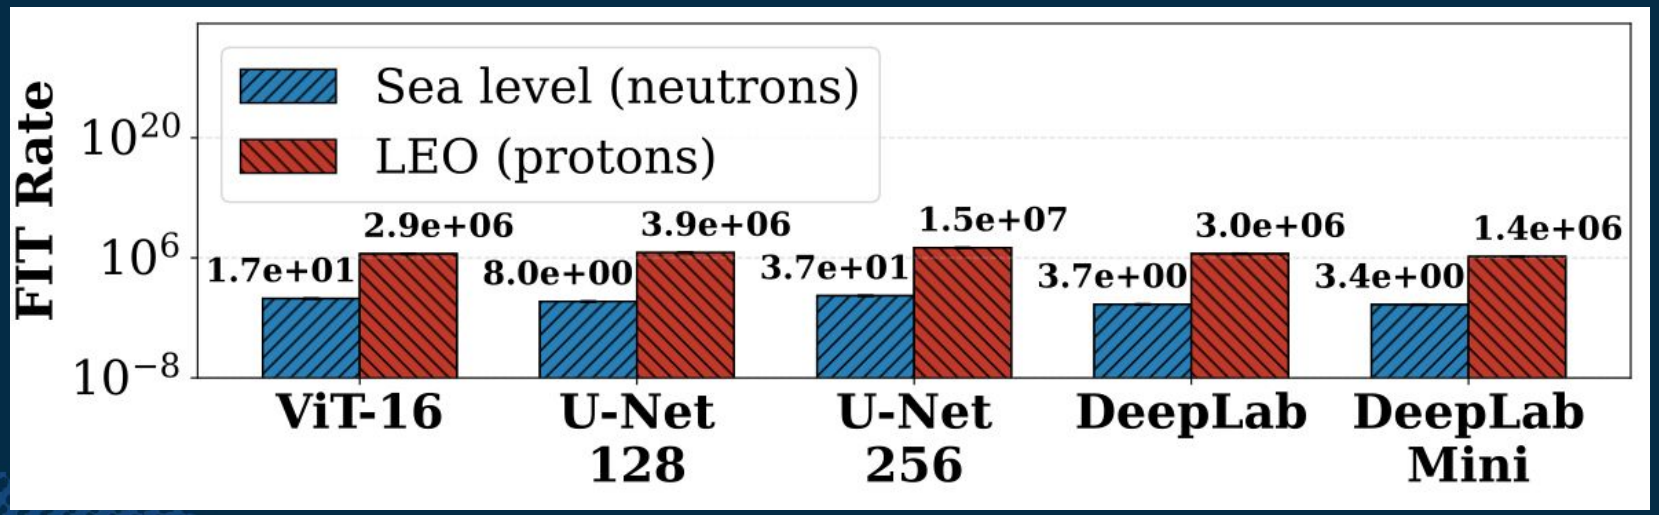
protons have higher probability of an output error

≥5 orders

higher critical-SDC FIT in LEO

4.84–17.22%

of SDCs become critical



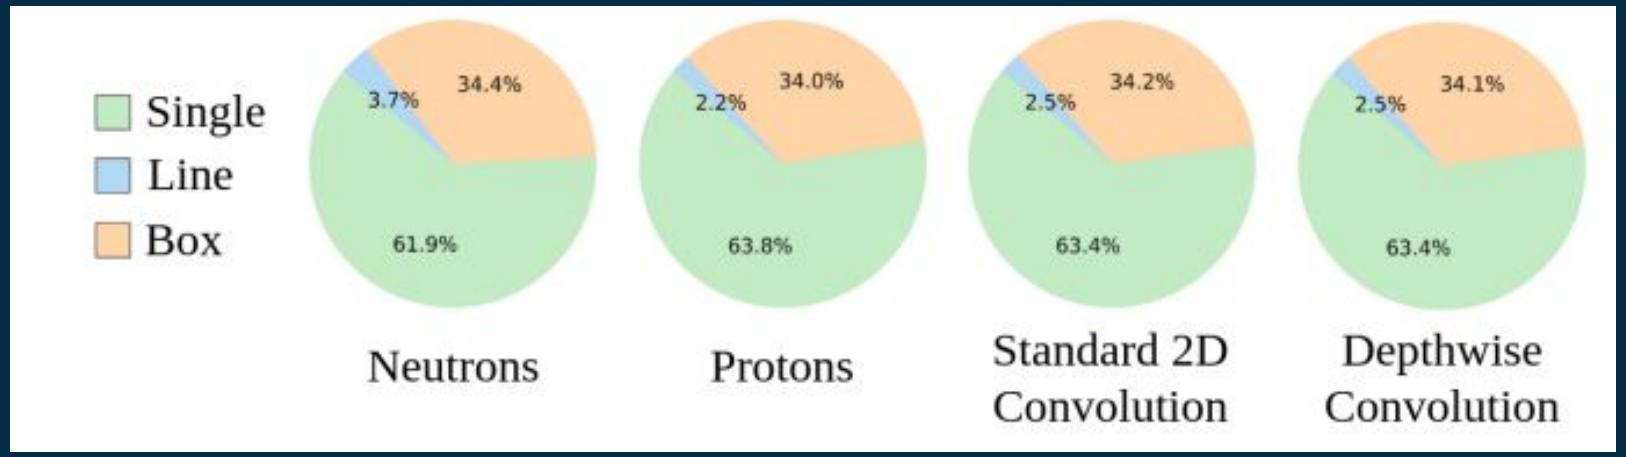
WHY THIS MATTERS

A raw SDC rate is not enough.

We need to know which errors are critical — and what hardware causes them.

The beam tells us the scale of the problem. It does not tell us the root cause.

The key discovery: radiation produces structured errors



SINGLE — 63.3%

One corrupted output location.

LINE — 2.7%

Multiple locations in one row/column.

BOX — 34.0%

A rectangular region of corrupted outputs.

90.8% of corrupted elements differ by only ±1 — but the spatial pattern determines propagation.

The pattern is intrinsic to the hardware

Across particles and convolution configurations:

92.03%
global similarity index

98.13%
spatial-distribution overlap

94.04%
error-delta overlap

83.91%
corruption-count overlap

FREQUENCY

Particle type + workload affect how often an error occurs.

MANIFESTATION

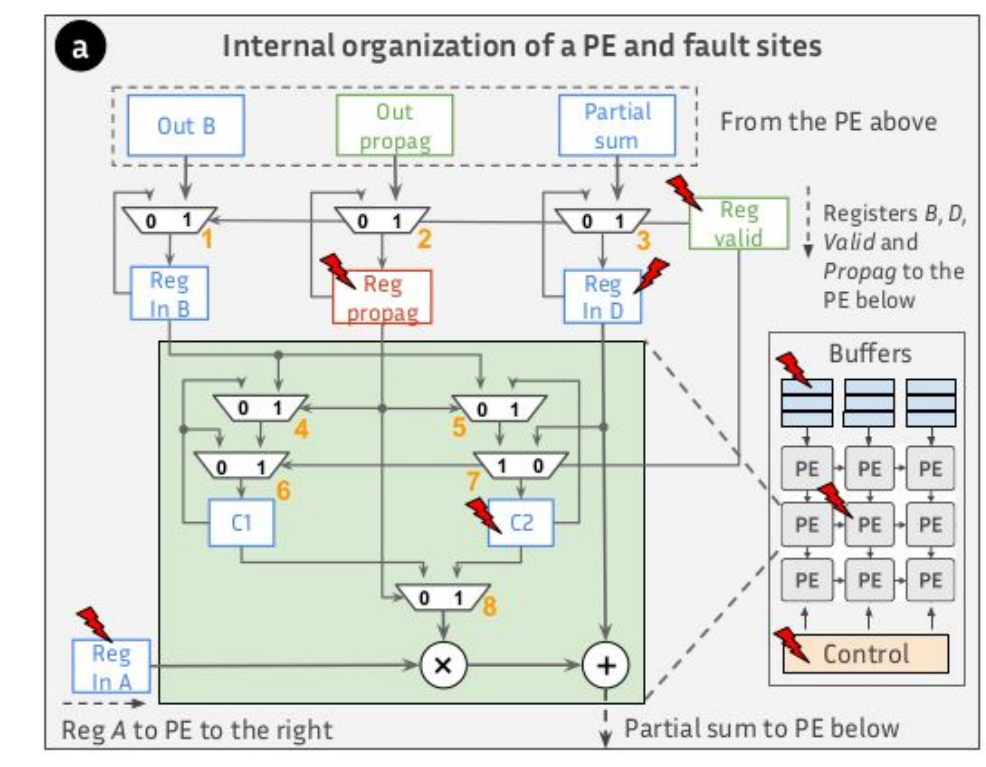
The systolic-array architecture determines the geometry of the corruption.

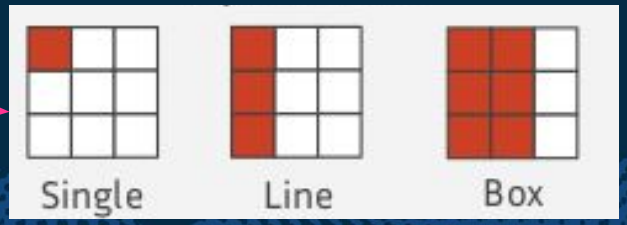
CONSEQUENCE

We can build a unified pattern-based fault model for software FI.

This is the crucial abstraction: model the shape of the error, not just a bit flip.

RTL + gate-level injection answers the question the beam cannot





SINGLE

MAC logic + partial-sum registers

LINE

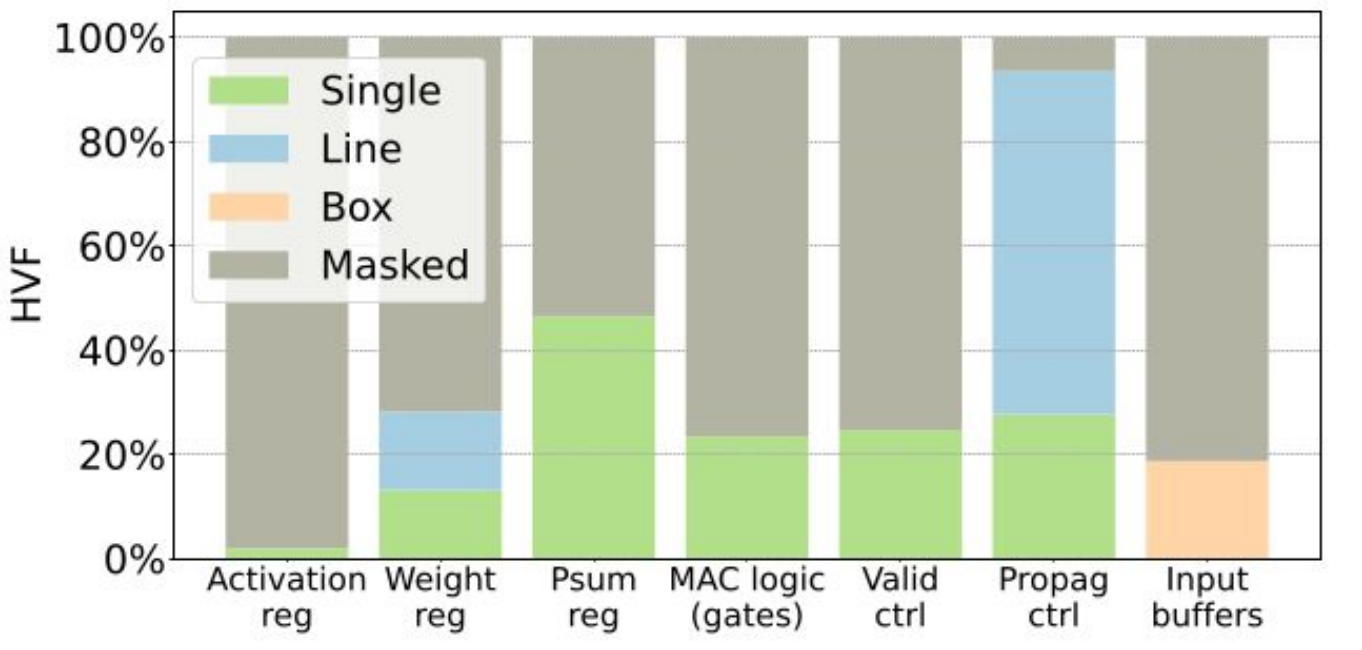
Weight registers + small control registers

BOX

Input buffers + im2col replication

30,000+ low-level injections
Fault location, cycle, bit: all recorded

RTL + gate-level injection answers the question the beam cannot



	Single	Line	Box
Radiation	63.4%	2.5%	34.2%
RTL/GL	58.4%	1.43%	40.1%
GE Area	52.0%	1.54%	49.9%

SINGLE
MAC logic + partial-sum registers

LINE
Weight registers + small control registers

BOX
Input buffers + im2col replication

30,000+ low-level injections
Fault location, cycle, bit: all recorded

Why a tiny hardware fault can become a large tensor error

1. PARTIAL SUM / MAC

Fault
↓
One PE's accumulated value
↓
One output location

→ SINGLE pattern

2. WEIGHT / CONTROL

Fault persists across cycles
↓
Weight or routing changes
↓
Multiple PEs / outputs

→ LINE pattern

3. INPUT BUFFER

One pixel is reused by im2col
↓
Appears in several patches
↓
Spatially repeated corruption

→ BOX pattern

The geometry is not accidental: it is a fingerprint of where the fault lives in the dataflow.

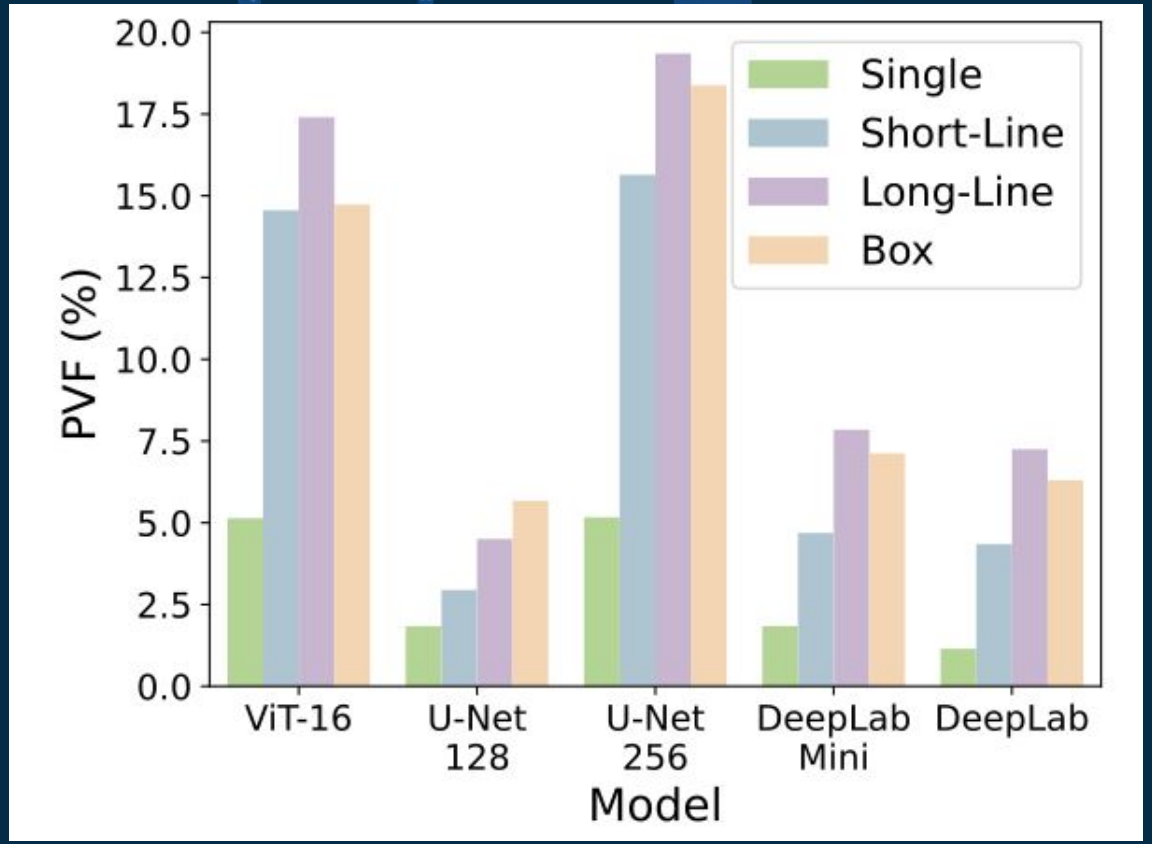
Software injection asks the final question: does the model care?

5.66M

software injections

12%

average injected faults → misprediction



Pattern sensitivity

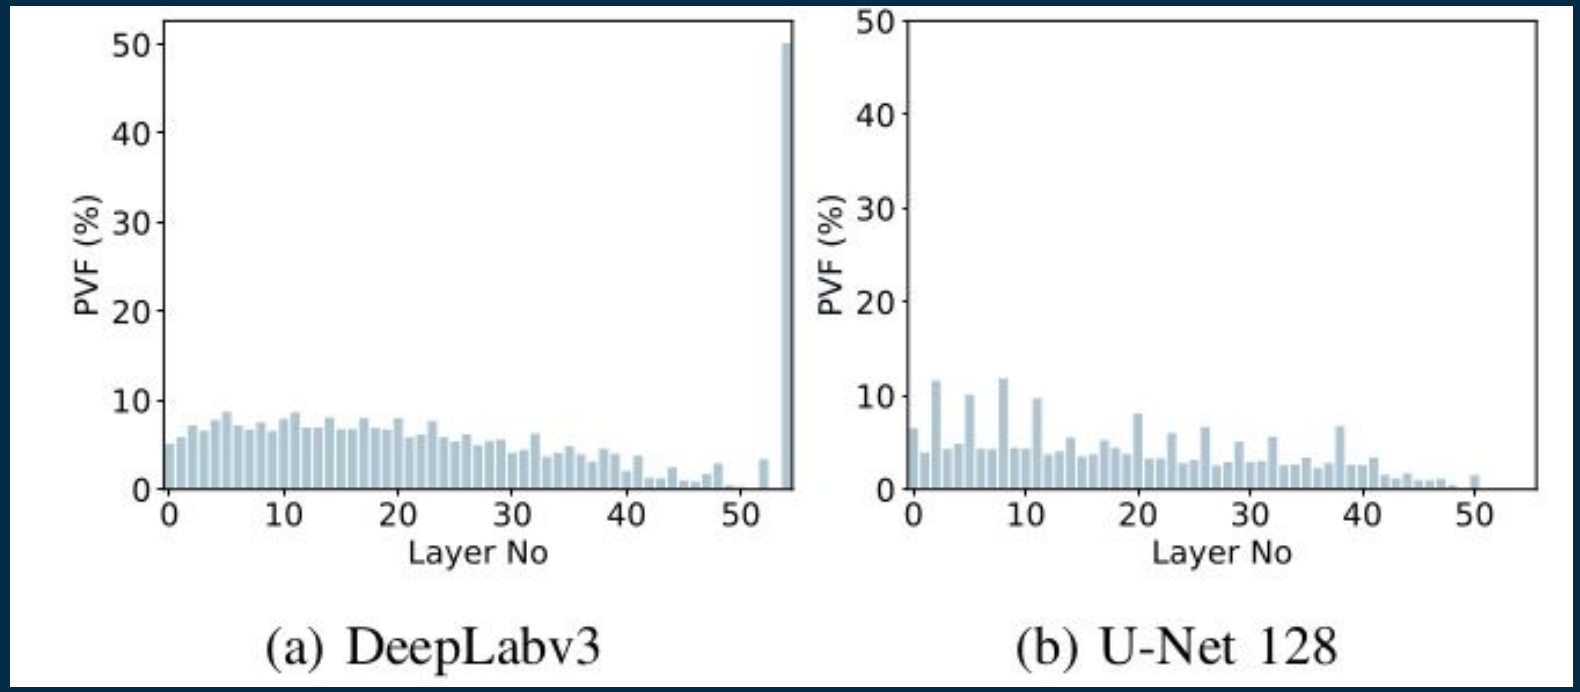
SINGLE

Low PVF — usually masked by the network.

LINE / BOX

High PVF; critical-SDC PVF can reach 19%.

And software vulnerability is not uniform across the model



DEEPER MODELS

Line faults can amplify across layers.

SHALLOWER MODELS

Box faults are more likely to remain critical.

U-NET

Skip connections can mask faults — but vulnerable encoder regions still emerge.

PVF gives a layer-by-layer map of where software is vulnerable.

Cross-layer AVF turns observations into a protection strategy

AVFpattern = HVFpattern × PVFpattern

SINGLE

High HVF
Low PVF
Large area: MACs ≈ 43% of vulnerable area

→ Favor software/algorithmic protection

LINE

Low HVF
High PVF

Generating resources <2% area

→ Protect control/pipeline registers

BOX

High HVF
High PVF

Input buffers ≈ 50% area

→ ECC is attractive, but incomplete

Targeted protection of key input buffers + registers could reduce critical fault probability by up to 87%.

Paper estimates: ECC alone ≤46% effective FIT reduction; control/pipeline replication ≈42% theoretical FIT reduction.

SC26

Main Conclusions

1. Radiation reliability is a cross-layer problem — a beam result alone cannot explain a misprediction.
2. Systolic arrays turn one transient fault into structured tensor errors: Single, Line, and Box.
3. The spatial pattern matters more than the simplistic “one bit flip” abstraction.
4. Cross-layer HVF × PVF identifies which hardware resources actually deserve protection.
5. The goal is not to harden everything — it is to harden the resources that create critical errors.