PINNeAPPle

STUDY REPORT · PUBLIC EXAMPLE

Ahmed body drag, mesh and what RANS cannot see

A 1984 benchmark, four meshes, 48 simulations and a surrogate — with model error separated from mesh error.

1.1%

of GCI on the fine mesh, observed order 4.26

+4.3% · −4.7%

of CD at 25°, on meshes L1 and L2

48 of 48

DOE cases run in OpenFOAM

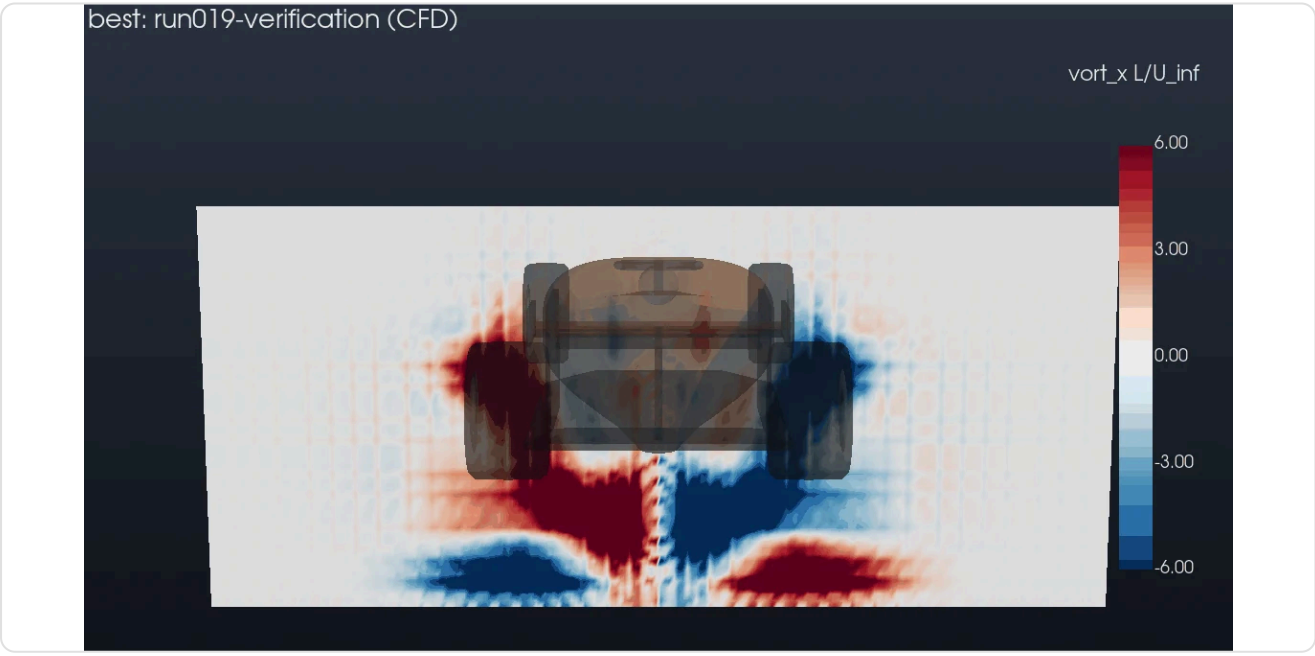

Document AER-01 · version 1.0 · 2026
Open case: Ahmed, Ramm & Faltin (1984)
OpenFOAM v1912 · gmsh 4.12 · half model, moving road, RANS k-ω SST
PINNeAPPle · LAIS & YAN LTDA · pinneapple-labs.github.io

The Ahmed body is the most used benchmark in automotive aerodynamics: a simple geometry with a slanted rear whose angle decides the drag. It was adopted because it is published, measured and hard — the flow separates, and that is where models fail.

ITEM	VALUE
Geometry	1.044 × 0.389 × 0.288 m · 0.1 m front radii · 0.222 m rear slant, stilts left out
Model	half model, moving road, steady RANS k-ω SST with wall functions, 5 prism layers
Variables	slant 0 to 40° · ground clearance 30 to 80 mm · speed 20 to 60 m/s
CAD	OpenCASCADE B-rep through gmsh, re-imported from its STL: closed, symmetric to 1.0 mm, volume 0.10993 m^3 against 0.10959 parametric

CFD on the CAD body at the reference condition (25°, 50 mm, 60 m/s) gave **CD 0.3004** against 0.2966 for the parametric body — 1.3% of difference from going through CAD and back. That is the study's first uncertainty, and it appears before the first simulation.

> **The question of the study.** It was not “what is the CD”: it was **how much of the error is mesh and how much is model** — because the two ask for opposite decisions. A mesh is refined; a model is replaced.

Four meshes, constant refinement ratio of $\sqrt{2}$, same reference condition, 1,000 iterations.

MESH	CELLS	CD	RUN CHECK
L0	84,223	0.2890	`rejected (mass balance)`
L1 (campaign)	207,080	0.2966	`review (residuals)`
L2	534,472	0.2707	`review (residuals)`
L3	1,398,073	0.2640	`review (residuals)`

L0–L1–L2 converges oscillatorily; L1–L2–L3 converges monotonically, with observed order **4.26** and Richardson value **0.2617**. The GCI on L3 is **1.1%**.

> **Read the order with suspicion.** 4.26 sits above the scheme's formal order (2), which makes the GCI a **lower bound** on numerical uncertainty — not a sharp estimate. Saying so is part of the result.

The campaign mesh (L1) sits **13.3% above** the Richardson value. It was kept for cost: 7.5 min against 19 min on L2 and 107 min on L3, on 4 cores. That 13% is the numerical uncertainty of every surrogate trained on the campaign.

The Ahmed, Ramm and Faltin (1984) experiment at 60 m/s, on the two measured meshes.

SLANT	MEASURED	L1	L2
0°	0.250	0.2693 (+7.7%)	0.2547 (+1.9%)
12.5°	0.230	0.2557 (+11.2%)	0.2399 (+4.3%)
25°	0.285	0.2973 (+4.3%)	0.2715 (−4.7%)
30°	0.378	0.3378 (−10.6%)	0.3184 (−15.8%)
35°	0.260	0.3835 (+47.5%)	0.4079 (+56.9%)
40°	0.250	0.3005 (+20.2%)	0.3958 (+58.3%)

Below the critical angle, refining the mesh brings the CFD towards the measurement — the behaviour expected from a numerical error. That band is what holds up the rest of the study.

> **Where validation stops holding.** From 30° up, refining makes it worse: the error grows from 10.6% to 15.8% and from 47.5% to 56.9%. That is not mesh noise — it is the model.

Steady RANS with k-ω SST keeps the slant flow attached beyond 30°, where the experiment separates. The result is the drag peak moved from 30° to 35–40°.

It is an error of **model physics**, and the practical consequence is direct: no amount of mesh fixes it, and no surrogate trained on those cases learns the real behaviour — it learns the solver's bias, faithfully.

> **What that costs a project.** Rear-geometry choices live exactly in the 25°–35° band. A study that only reports “CD = 0.32” there is handing over a number with up to 57% model error — without saying so.

Forty-eight geometries by maximin Latin hypercube (centred L2 discrepancy 0.0039 against 0.0141 for random designs of the same size); 48 of 48 ran. The final set has 54 cases: 28 train, 6 validation, 10 test and 10 extrapolation (every case with slant above 34°).

TARGET	SELECTED FAMILY	TEST ERROR	TEST COVERAGE	EXTRAPOLATION ERROR
CD	GP	0.42%	1.0	5.2%
CL	GP	2.0%	0.9	18.4%
Cp (field)	POD-GP	9.9%	0.89	20.0%

Relative error for CD and CL; mean relative L2 for the Cp field. The test and extrapolation splits are measured once, after model selection.

Sobol indices over the training region: on CD, slant explains **0.96** of the variance, clearance 0.05 and speed 0.01. For friction drag the order flips (0.09 / 0.04 / 0.93): two phenomena, two dependencies.

> **Extrapolation warns.** Above 34° of slant, interval coverage drops to 0.6–0.8: the model is not to be used there, and the twin says so on screen instead of answering.

1.1%
of GCI on mesh L3

13.3%
of campaign-mesh uncertainty

57%
of model error at 35°

0.42%
of surrogate error on CD, test

What this study does not solve

- Steady RANS does not capture the unstable separation of the slant above 30°; that needs URANS or LES.
- The campaign-mesh uncertainty (13.3%) is embedded in every surrogate trained on it.
- The friction model fails the physical check “falls with Re” and was dropped — not for error, for physics.
- y^+ inside the wall-function range was required in every run; outside it, the result is not comparable.

Provenance

WHAT	WHERE
Reference experiment	Ahmed, Ramm & Faltin (1984)
Full case	`PINNeAPPle/examples/end_to_end/ahmed_body`
Data and evidence	`results/dataset` · `results/evidence.zip` (with the SHA-256 of every source)
Public article	`pinneapple-labs.github.io/artigos/aerodinamica-autonoma/`

This is an example report, on a public benchmark. A study for your body follows the same structure: a mesh with GCI, validation against an experiment, model error separated from numerical error — and the page of what is not proven.