ictive Modeling of G-Induced Loss of Consciousness (GLOC) Using Physiological Data and Machine Learning

Lucas R. Owen · Joshua B. Smith · Eunjun Shin · Sarah Bolton

Department of Computer and Cyber Sciences, United States Air Force Academy

2026 IEEE National Aerospace and Electronics Conference
August 9-12, 2026

Paper ID: 47

The Problem: A Hazard Caught Too Late

GLOC

Sustained G-forces cut cerebral blood flow, causing temporary incapacitation. Even brief loss of control can be fatal (CFIT).

CFIT

Disorientation or unconsciousness at low altitude leads to Controlled Flight Into Terrain.

Auto-GCAS is reactive

It intervenes only once ground impact is imminent, and cannot adapt to the pilot's physiological state.

AUTO-GCAS, SINCE 2014

13

pilots saved

12

aircraft saved

...but only after impact was imminent.

We want to see it coming: a predictive layer that reads pilot state could act seconds before incapacitation, not after.

Research Question and Contribution

Binary classification: **does this timestep fall within 10 seconds before a GLOC event?**

Leakage-free

We strip out every feature that reflects the centrifuge, not the pilot.

Subject-independent

Leave-one-subject-out evaluation, so the score reflects generalization to unseen pilots.

Cockpit-deployable

The winning signals come from joystick tracking and pupillometry, hardware you could field.

The Dataset: AFRL 711th Human Performance Wing

~1.4M ROWS

61 SUBJECTS

312 COLUMNS

- EEG 32ch
- Joystick
- Vitals
- Eye / pupil
- fNIRS
- Math task

NON-AFE (no anti-G augmentation)

34 / 34

ROR runs resulted in GLOC

*Every unprotected pilot lost consciousness. The task is **when**, not **whether**.*

AFE (anti-G suit + augmented flow)

5 / 27

ROR runs resulted in GLOC

Protective gear changes the outcome so strongly it becomes a confound (more soon).

Protocol: gradual onset at 0.1 G/s for baseline, then rapid onset at 3.0 G/s while performing the AGSM, until GLOC.

The Instrumentation

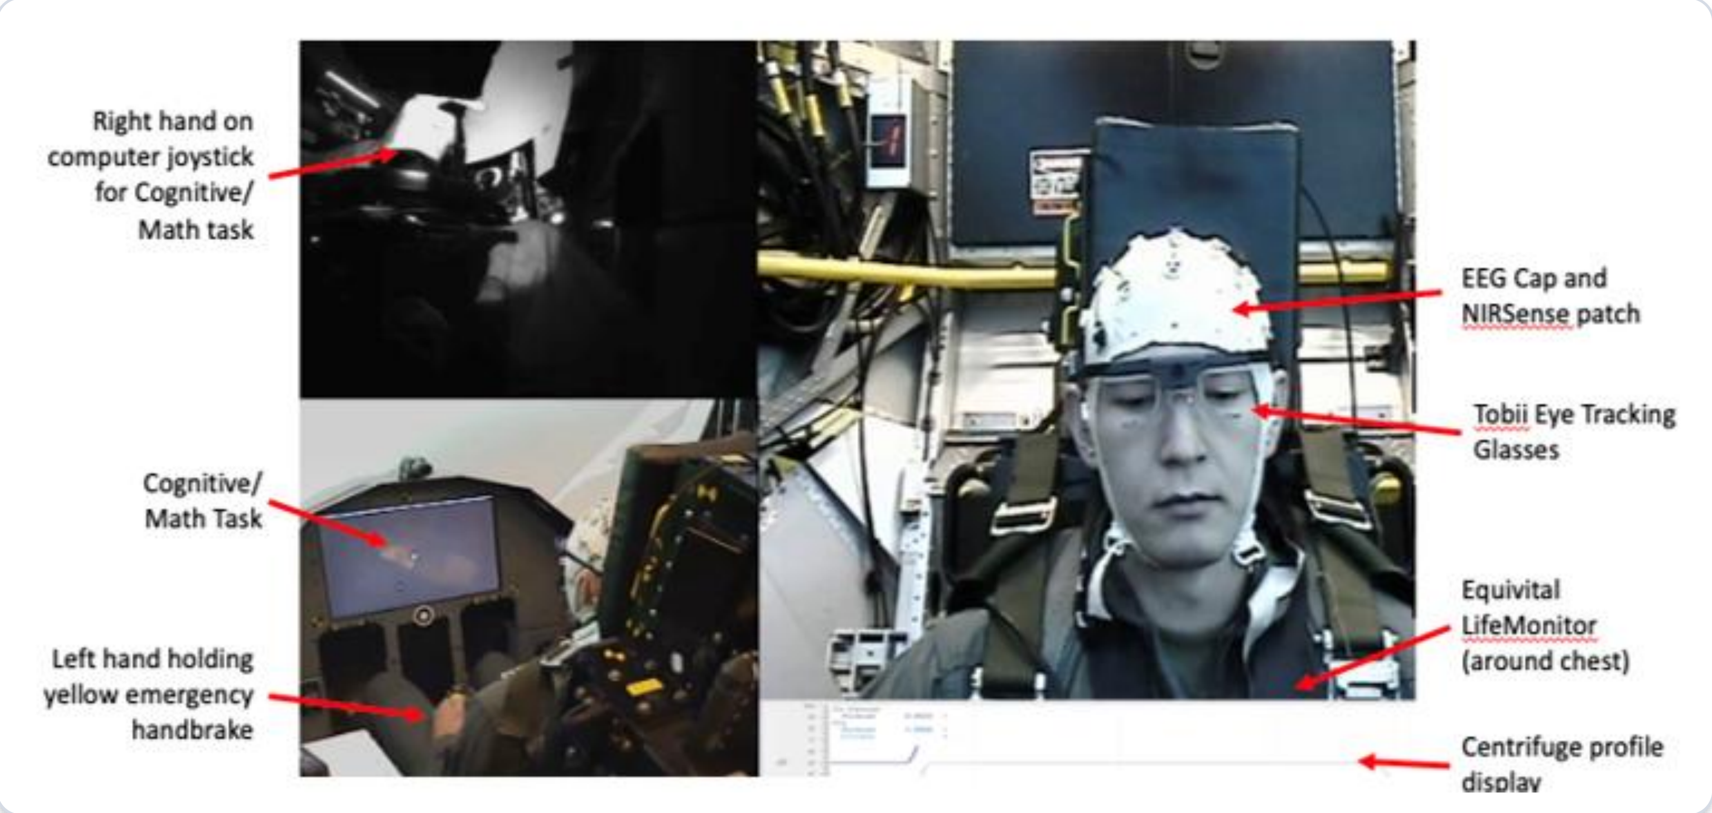

Six sensor systems on one subject. The two that carry the signal, joystick and eye tracking, are the two realistically fieldable in a cockpit today.

The Core Problem: Data Leakage

The model kept "cheating": four times the top predictor was the environment, not the pilot.

1. Centrifuge telemetry

Models read the programmed G-profile, the script of the run, not the pilot's response.

2. Device accelerometers

Top features were EEG-headset and eye-tracker IMUs: physical G on the devices, not brain or eye.

3. EEG motion artifacts

Parietal delta power spiked at or after onset: a textbook artifact from head slump.

4. Condition confound

GLOC was near-perfectly predicted by AFE vs Non-AFE. The model learns the label, not the pilot.

Fix: exclude all four, then restrict analysis to Non-AFE rapid-onset windows.

Pipeline: From Raw Data to an Honest Score

Raw data
61 subj · 312 cols

→

Restrict
Non-AFE ROR window

→

De-leak
remove 4 sources

Clean set
22–24 features

→

Random Forest
100 trees, 3:1

→

LOSO-CV
34 held-out pilots

Every arrow discards information that looked useful but was not the pilot. What survives is small and defensible.

Model and Honest Evaluation

Random Forest, 100 trees

- Robust out of the box, minimal tuning
- Feature importances for interpretability
- Tolerates mixed feature types
- 3:1 undersampling of negatives for balance

XGBoost matched it but needed far more tuning; KNN and SVM did not scale.

Leave-one-subject-out CV

Train on every pilot but one. Test on the pilot the model has never seen. Repeat for all 34.

The only method that answers the real question: will it work on the next pilot, not the next timestep?

We report **Recall / Precision / F1**, not accuracy: the class imbalance makes accuracy meaningless.

RESULTS

The journey, one row per methodological change

Configuration	Recall	Prec.	F1
Handpicked features (30)	0.82	0.46	0.59
All 213 features	0.63	0.50	0.56
No centrifuge columns (198)	0.66	0.33	0.44
Pure pilot-state (168)	0.62	0.21	0.31
Temporal XGBoost (540)	0.52	0.26	0.35
Temporal RF 20 s (162)	0.45	0.29	0.35
No EEG / no accelerometers (24)	0.39	0.16	0.22
Non-AFE ROR only, no fNIRS (22)	0.86	0.67	0.75
Non-AFE ROR only, with fNIRS (24)	**0.87**	**0.72**	**0.78**
fNIRS subjects only (24)	0.80	0.68	0.73

Notice the dip in the middle: stripping leakage first makes the score look worse. That is the model getting honest before it gets good.

Results

Best configuration – Non-AFE rapid-onset windows, leave-one-subject-out:

0.86
RECALL
86% of pre-GLOC timesteps caught

0.67–0.72
PRECISION
without / with fNIRS

0.75–0.78
F1
harmonic mean

Removing leakage first makes the scores look ***worse*** before they recover. That is the model getting honest before it gets good, and the dip is the evidence the final number is real.

fNIRS Is Strongest – But Not Required

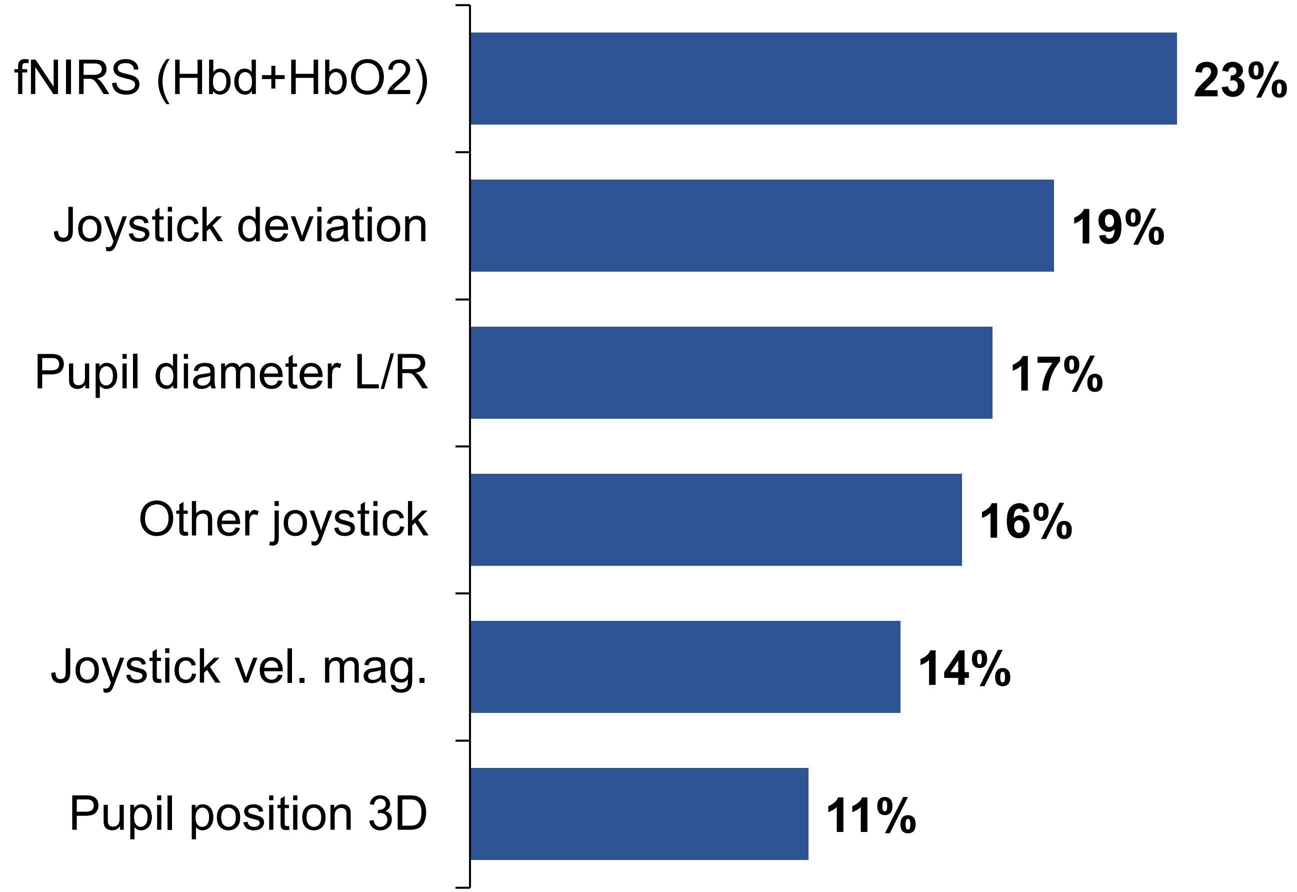

fNIRS = the mechanism

Top signal (~23%); directly measures the brain hypoxia that causes GLOC.

...but remove it and:

Recall 0.87 → 0.86

Precision 0.72 → 0.67

Joystick + pupillometry alone are deployable today.

Why EEG Failed; Why Condition Rules

EEG delta power around GLOC onset

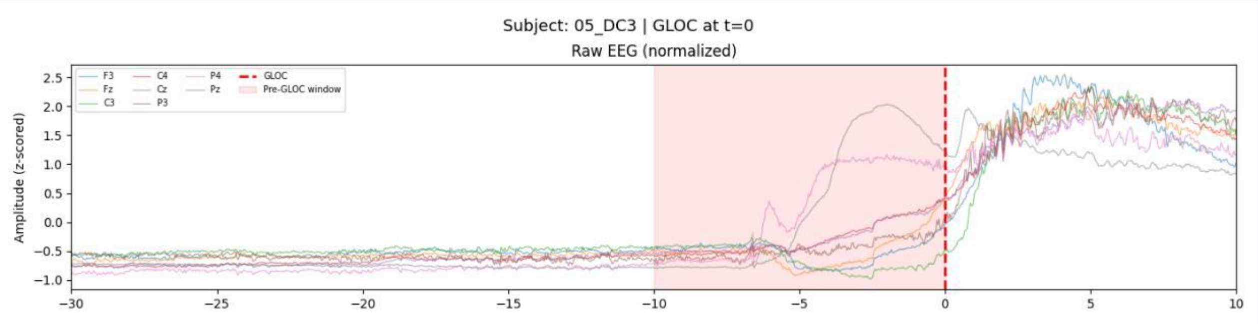

The spike lands **at or after** onset, never before: the head slumping at loss of muscle tone. A motion artifact, not a warning.

Condition is the real driver

AFE vs Non-AFE predicts GLOC almost perfectly on its own.

Any model spanning both conditions learns the label, not the pilot.

Restricting to Non-AFE is what makes the evaluation honest.

Operational Deployment

Prompt AGSM

~3 s is enough to cue or reinforce the anti-G straining maneuver.

Pre-arm Auto-GCAS

1–2 s of anticipation lets recovery logic engage with less delay.

Hand off to autonomy

On high-confidence predicted incapacitation, autonomy holds the jet until recovery.

Cost asymmetry favors recall: a nuisance alarm is cheap, a missed GLOC is not.

Limitations and Future Work

LIMITATIONS

Small positive class
Only 39 GLOC events across 61 subjects. Interpret with caution.

fNIRS reliability
35% of subjects lost fNIRS to sensor disconnection.

EEG unused
Artifact rejection tuned to high-G could recover real neural signal.

FUTURE WORK

Shorter lookahead
The 10 s window was not optimized; test 5 s for tighter precision.

Response time as label
More sensitive than accuracy; a path to recovery prediction.

Larger, varied pools
Validate across studies and demographics before any fielding.

Summary

GLOC onset is predictable with cockpit-deployable sensors, once the leakage is gone.

0.86	0.67–0.72	0.75–0.78
RECALL	PRECISION	F1

Joystick tracking and pupillometry carry the signal; fNIRS helps but is not required.

Lesson: in physiological ML, most of the work is proving the model isn't cheating.