

AI-Assisted Survival Heatmap


for Tactical Decision Support
in Mountainous Operations

Jung-Tae Lee, Eunae Lee and Muhyun Byun

Defense AI R&D Institute

Agency for Defense Development, Republic of Korea

Why mountain operation




The research problem

Spatial prediction


How does survival vary across a terrain and force configuration?

Decision alignment

How should a shortlist reflect the mission's objective?



A research framework



Prediction layer


A dense survival field learned from simulated trajectories

Decision layer

A ranked shortlist conditioned on explicit mission criteria

1. Combat simulator

Aggregate Lanchester models are too coarse; JCATS / OneSAF / SWORD are too costly to instrument per cell.



Left: phase controller. Right: the six operations executed every $\Delta t = 30$ s tick.

≈1.5 kLoC of Python

Scripted discrete-event loop; full observability of per-tick, per-soldier state at $\approx 1.5 \times 10^5$ ticks·s⁻¹.

Slope

Five-tier speed multiplier from military movement-rate tables: 1.00 below 8°, down to 0.30 above 35°.

Weather

Sampled in 2–6 h blocks so conditions persist: clear 35 %, cloudy 25 %, rain 22 %, fog 10 %, storm 8 %. Fog compresses detection far more than movement.

Soldiers

Ten per team; base speed 0.55–1.20 m·s⁻¹, detection 200–1000 m. Stamina 0–100 steps speed down to 0.40.

2. Training simulation data

Data, labels and training

9,999

independent runs generated

7,814

engagements kept for training

850

scenarios held out

16 h

on a single A100

Labels come from the paths soldiers actually walked

- Every cell traversed by a BLUE soldier who lived is labelled 1; every cell traversed by one who died is labelled 0
- A binary mask $M \in \{0,1\}^{512 \times 512}$ records visited cells, so the loss scores only labelled ones
- Given a snapshot minutes before contact, the network predicts the probability that a friendly soldier occupying a cell survives


Why we do not report IoU or Dice

- Standard segmentation metrics are scored only on occupied cells, so they say nothing about the dense field the system actually uses
- Training: AdamW, $\text{lr} = 1 \times 10^{-4}$, weight decay 1×10^{-4} , batch 32, 150 epochs, mixed precision
- Operational utility is measured instead — by Monte-Carlo replay (slide 11)


2. Training simulation data

X = nine masked channels at the snapshot, taken about ten minutes before contact. Y = the fate of every cell the BLUE entities then walked. 8,700 such windows were extracted as training pairs.

X — input at snapshot time




Y — label over the engagement window



2. Training simulation data

Nine semantic input channels (csv -> image)



Channel 0

Normalized elevation

Channels 1-2

BLUE and RED occupancy


Channels 3-8

Weapon-class maps

Three classes per team

2. Training simulation data

Swin-Tiny encoder with a CNN decoder



Architecture

- Input stem widened to nine channels; feature maps at 128^2 , 64^2 , 32^2 , 16^2 with 96 / 192 / 384 / 768 channels
- Decoder mirrors it with bilinear $\times 2$ upsampling and skip connections; each stage concatenates the matching encoder map, then two 3×3 Conv-GELU blocks
- A final 1×1 convolution returns logits at input resolution; masked loss scores visited cells only


33.75 M

parameters

< 1 s

inference on one T4

3. Prediction heatmap



Probability map

Green

to be better

Red

to be worse

Which one would you choose?

There are too many green zones.

A probability field is not a decision.

4. Recommends Top-k

1 Feasibility screening


Within 600 m, with at least 150 m spacing
Mean patch survival ≥ 0.85

2 Candidate characterization

Counterfactual, geometric and route-based evidence

3 Mission-conditioned scoring

Thirteen criteria, normalized within the candidate set
Weighted additive score conditioned on the mission



5. Validates Top-k

Replay tests the effect of the recommendation


Policy	Role in the comparison
Baseline	Default simulator phase controller
Random feasible position	Effect of repositioning within the feasible region
AI Top-1, Top-2, Top-3	Effect of selecting recommended positions

850 scenarios × 30 replays × 5 policies

127,500 simulated engagements

5. Validates Top-k

Scenario means over 850 held-out scenarios, 95 % t confidence intervals.



Rank	Mean predicted probability	Replay BLUE survival
Top-1	0.960	0.720
Top-2	0.920	0.711
Top-3	0.890	0.698


Calibration

Results are obtained without model calibration.

Random arm

BLUE movement reduces RED firing opportunities


5. Validates Top-k



Threat avoidance: 5, 6, 9 Fast, safe movement: 3, 4, 8

5. Validates Top-k

Normalised Mission Outcome (NMO): every one of the twelve candidates replayed 30 times, scored on the metric central to each mission.



Clearer separation

High ground
Fast, safe route

Weak alignment

Mutual support

Future work
Ranking stability

Conclusion

- Trained model-guided repositioning improves survival rates

0.558 **baseline**, 0.720 **Top-1**

- Decision support layer should be tailored to mission requirements

Limitation

- ❖ **System-in-the-loop**
: the network is trained on simulator ground truth and validated by re-injecting recommendations into the same simulator
- ❖ **Trained data**
: Limited geographic data constrain generalizability; real-world KCTC data could enhance practical applicability

Thank you

E-mail: jtlee87@add.re.kr
jtlee87@naver.com

Backup

Simulator detail, held for questions.

Backup — environment model: terrain, weather, stamina

All three scale the same underlying quantities: how fast you move, how far you see, how often you hit.

Slope-based speed multiplier

Slope	Speed ×
-------	---------

< 8°	1.00
------	------

8–15°	0.85
-------	------

15–25°	0.68
--------	------

25–35°	0.50
--------	------

≥ 35°	0.30
-------	------

Weather — sampled in 2–6 h blocks, so it persists across an engagement

Weather	Detection	Speed	Accuracy	Prior
Clear	100 %	100 %	100 %	35 %
Cloudy	85 %	95 %	90 %	25 %
Rain	70 %	75 %	75 %	22 %
Fog	40 %	82 %	70 %	10 %
Storm	45 %	55 %	55 %	8 %

Stamina (0–100)

- ≥ 60 nominal ×1.00; 40–60 ×0.80; 20–40 ×0.60; < 20 ×0.40 plus a forced halt
- Rest recovers at 4.0 per tick against 0.8 spent moving — five times faster
- Exhaustion is a positive feedback loop: slower movement raises exposure, which raises damage

Fog is the decisive variable

Fog cuts detection to 40 % while barely slowing movement, creating an asymmetric state in which a squad can reposition but cannot see — which is exactly when position selection matters most. Storm, by contrast, halves everything and simply suppresses engagement on both sides.

Backup — force composition and the phase state machine

Ten entities per team: 2 reconnaissance, 7 riflemen, 1 sniper — each running a different tactical behaviour.

Per-role parameters

Role	Base speed	Detection	Max range
Recon (2)	1.20 m/s	200 m	300 m
Rifleman (7)	0.90 m/s	200 m	300 m
Sniper (1)	0.55 m/s	1000 m	1000 m

The sniper is deliberately given low mobility so it holds high ground on its own, and it is excluded from the team centroid used for objective computation.

Detection

Binary test on $\text{range} \leq \text{detection} \times \text{weather visibility}$. Symmetric ($A \leftrightarrow B$); effective range falls linearly as weather degrades.

Engagement

Hit probability = band accuracy $\times$ weather accuracy. Damage 25–55 HP, so one to three hits are lethal. Snipers prefer the farthest target, everyone else the nearest.

Five-state phase controller



Transitions: one hour elapsed (infiltrate $\rightarrow$ search) $\cdot$ enemy detected ($\rightarrow$ engage) $\cdot$ friendly losses above 50 % ($\rightarrow$ withdraw).

criteria	THREAT_AVOID	FAST_SAFE_ROUTE	HIGH_GROUND	SUPPORT	SURVIVAL_MAX	BALANCED
p_stay	.10	.10	.17	.15	.55	.20
worst_unit_threat	.25	.10	.15	.10	.15	.13
threat_red	.10	—	—	—	—	—
threat_SR	.12	—	—	—	—	—
threat_HE	.08	—	—	—	—	—
exposure	.20	.10	.10	—	.10	.10
eta_min	—	.30	—	.10	.05	.13
path_risk	.10	.25	—	.10	.10	.12
path_min_p	—	.15	—	—	—	.05
elev_adv_m	.05	—	.28	—	—	.10
terrain_inf	—	—	.20	—	—	.05
support_inf	—	—	.10	.32	.05	.07
dist_from_ally_m	—	—	—	.23	—	.05
weight_sum	1.00	1.00	1.00	1.00	1.00	1.00