NEURIPS 2026 · STATIC DEMO

EchMo

End-to-End Radar Echo Segmentation

The first end-to-end semantic segmentation architecture for pulse-Doppler radar.

Echo in, mask out.

Liwen Zhang[†*] · Xinying Fu[†] · Youcheng Zhang[†] · Zijun Hu · Wen Chen · Shi Peng · Jie Zhou · Zhe Ma[*]

[†] Equal contribution · [*] Corresponding authors

Intelligent Sci. and Tech. Academy (ISTA), CASIC, Beijing 100041, China

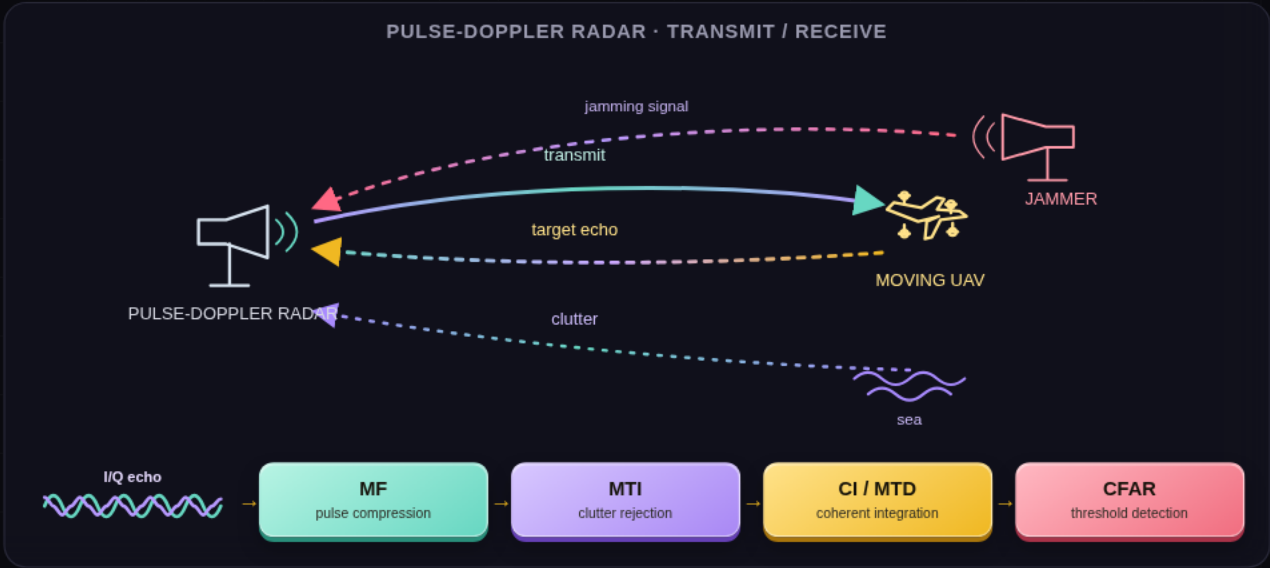
Perception & Cognition Laboratory (PeCoLab@ISTA), Beijing, China

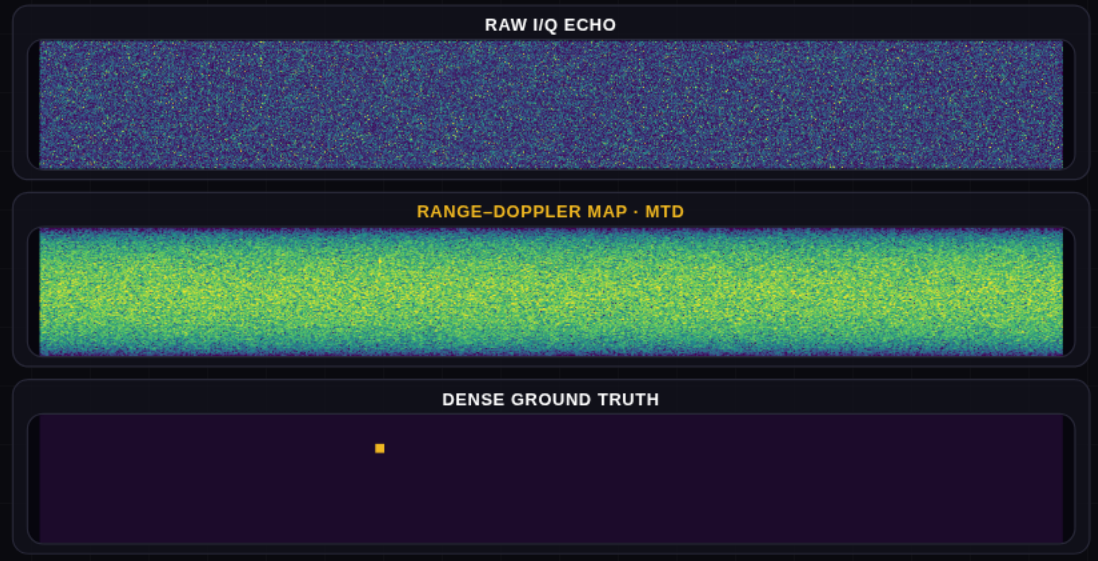
Physics-initialized learnable RSP | ReCK — how much prior to keep? | 1.29M params · 97.6 mIoU

THE PROBLEM

Radar observation for target perception, and the bottleneck of classical signal processing.

A pulse-Doppler radar illuminates the scene with coherent pulse trains; the receiver observes complex I/Q samples in which **faint target echoes, ground/sea clutter, noise, and active jamming** are superimposed along fast and slow time. Classical radar signal processing (RSP) squeezes out range–Doppler information through a fixed chain of matched filtering, MTI, coherent integration, and CFAR detection.





Idealized assumptions
Classical operators depend on priors such as noise and clutter models.

Shallow, isolated optimization
Stages are tuned separately, without end-to-end coordination.

Information loss in stages
Converting raw echoes to intermediate representations may discard detail.

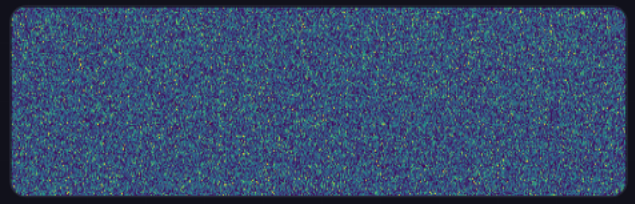
Can the complete mapping—echo in, semantic mask out—be learned without optimizing each stage in isolation?

DEEP LEARNING ON INTERMEDIATE REPRESENTATIONS

The deep learning paradigm built on conventional intermediate representations

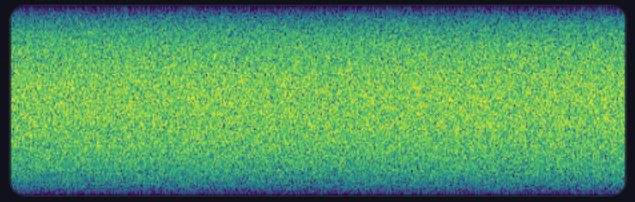
Mainstream radar-perception deep learning feeds intermediate representations produced by conventional radar signal processing (RDM, RAD tensors) into networks for target detection and semantic discrimination. These methods learn only the **"representation → task"** mapping; the transform from raw I/Q echoes to that representation still relies on a handcrafted signal-processing pipeline, so the inherent information loss is locked in before the network input.

RADAR ECHO · PHYSICAL SAMPLES



fixed RSP front end →

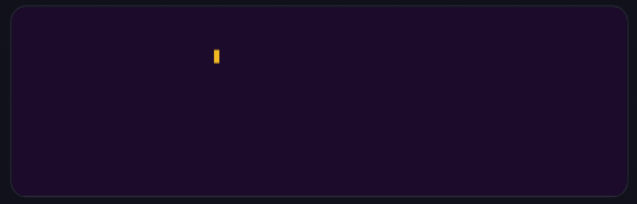
DEEP MODEL INPUT · INTERMEDIATE REPRESENTATION



RDM · RAD · ...

task mapping →

DENSE SEMANTIC SUPERVISION



TWO MAINSTREAM TASKS · REPRESENTATION → SEMANTICS

RD / RAD
RSP representation

→ **CNN / Transformer**
generic / radar-tailored

→ **ROD · radar object detection**



output · sparse boxes

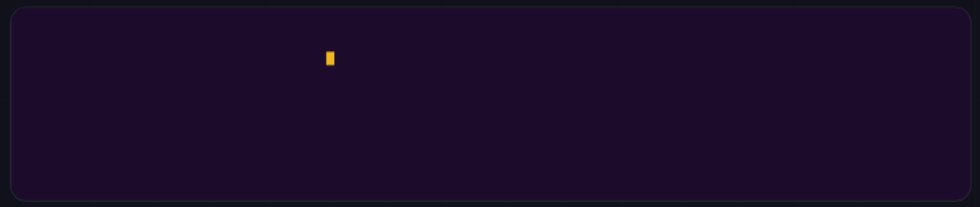
→ **RSS · radar semantic segmentation**



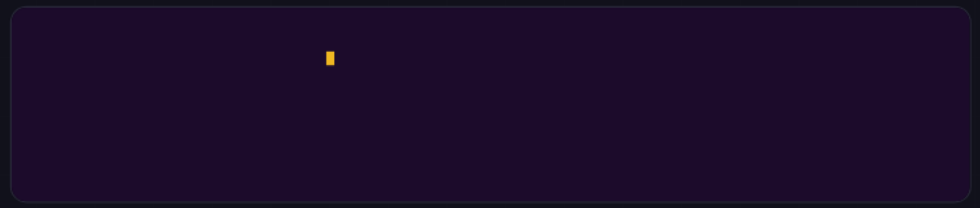
output · dense semantic mask

Radar-tailored models consistently outperform generic vision backbones; yet the fixed RSP front end causes irreversible information loss, and echo-to-semantics end-to-end optimization remains impossible.

DeepLabv3+ · CNN



ViT · Transformer



On simple high-SNR tasks, mainstream segmenters achieve good masks once conventional processing has extracted the features.

Deep models learn semantic mapping, but the physical processing before input reaches the network remains outside joint learning.

THE METHOD · REINTERPRETATION

From raw echoes to semantic masks: the EchMo architecture

Physical priors are no longer fixed code: they become initialization and differentiable computation flow, adapted by data within a physics-shaped architecture.

input: B×C×P×R complex I/Q tensor

▶ ANIMATE ARCHITECTURE

Forward signal flow in both pipelines
Classical operator reinterpreted as a learnable module

CLASSICAL · MF
Matched filtering
Pulse compression · SCNR gain

CLASSICAL · MTI
Delay-line canceller
Stationary-clutter rejection · Doppler high-pass

CLASSICAL · CI / MTD
Fourier Doppler filtering
Slow-time integration · range-Doppler map

CLASSICAL · CFAR
Detection & decision
Threshold / local background

Reference Init · conjugate-reversed transmit waveform
LDC Init · binomial delay taps
DF Init · Fourier cosine/sine kernels
CFAR-inspired · local peak-contrast kernel

ECHMO · 01
LMF
K parallel chains · reference-seeded complex conv

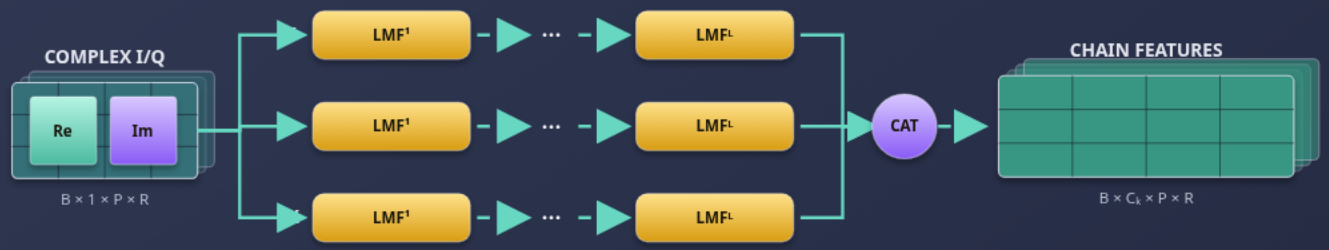
ECHMO · 02
MLDC
N+1 parallel orders · padding · delay cancellation · channel concat

ECHMO · 03
GDF
Grouped Doppler filters · cos/sin init · learnable offsets

ECHMO · 04
PKC-Seg
PeakConv encoder → decoder → dense mask

LMF · Learnable matched filter

Reference-seeded complex convolution; K parallel Chains concatenate into feature channels.



$$K_{nc}^{\dagger}\{X;K^{\dagger}\} = K_{nc}\{\mathrm{Re}(X);K^{\dagger}[1]\} - K_{nc}\{\mathrm{Im}(X);K^{\dagger}[2]\} + j\left[K_{nc}\{\mathrm{Re}(X);K^{\dagger}[2]\} + K_{nc}\{\mathrm{Im}(X);K^{\dagger}[1]\}\right]$$

REFERENCE INIT

$$K^{\dagger}[1] \leftarrow \mathrm{Re}(S^{*} \text{ or } \tilde{S}^{*}), \quad K^{\dagger}[2] \leftarrow \mathrm{Im}(S^{*} \text{ or } \tilde{S}^{*})$$

THE METHOD · REINTERPRETATION

From raw echoes to semantic masks: the EchMo architecture

Physical priors are no longer fixed code: they become initialization and differentiable computation flow, adapted by data within a physics-shaped architecture.

input: B×C×P×R complex I/Q tensor

▶ ANIMATE ARCHITECTURE

Forward signal flow in both pipelines
Classical operator reinterpreted as a learnable module

CLASSICAL · MF
Matched filtering
Pulse compression · SCNR gain

CLASSICAL · MTI
Delay-line canceller
Stationary-clutter rejection · Doppler high-pass

CLASSICAL · CI / MTD
Fourier Doppler filtering
Slow-time integration · range-Doppler map

CLASSICAL · CFAR
Detection & decision
Threshold / local background

Reference Init · conjugate-reversed transmit waveform
LDC Init · binomial delay taps
DF Init · Fourier cosine/sine kernels
CFAR-inspired · local peak-contrast kernel

ECHMO · 01
LMF
K parallel chains · reference-seeded complex conv

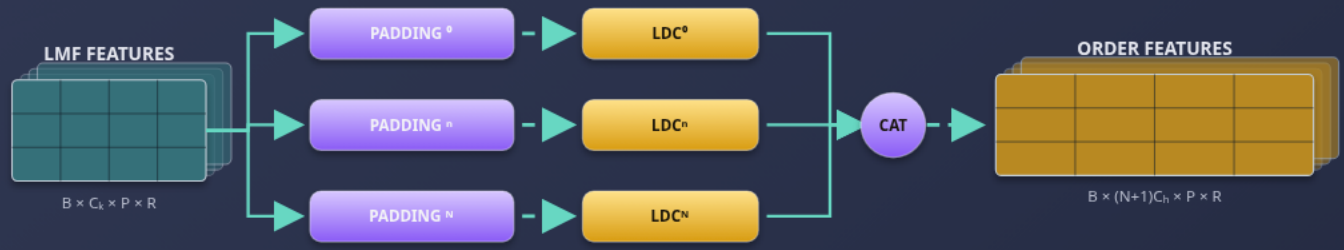
ECHMO · 02
MLDC
N+1 parallel orders · padding · delay cancellation · channel concat

ECHMO · 03
GDF
Grouped Doppler filters · cos/sin init · learnable offsets

ECHMO · 04
PKC-Seg
PeakConv encoder → decoder → dense mask

MLDC · Multi-order learnable delay canceller

N+1 parallel padded delay-canceller orders; outputs concatenate along channels.



$$\mathcal{H}^{*}_{caus}\left(\{\mathrm{Pad}^{n}_{caus}(\hat{X})\}_{n=0\ldots N};\ \{H^{n}\}_{n=0\ldots N}\right) = \tilde{X}^{0} \oplus \cdots \oplus \tilde{X}^{N}$$

LDC INIT

$$H^{n}[:,i] \leftarrow h^{n}_{i}\cdot\mathbf{1};\quad h^{n}_{0}=1,\ h^{n}_{i}=(-1)^{i}\binom{n}{i}$$

THE METHOD · REINTERPRETATION

From raw echoes to semantic masks: the EchMo architecture

Physical priors are no longer fixed code: they become initialization and differentiable computation flow, adapted by data within a physics-shaped architecture.

input: B×C×P×R complex I/Q tensor

▶ ANIMATE ARCHITECTURE

Forward signal flow in both pipelines
Classical operator reinterpreted as a learnable module

CLASSICAL · MF
Matched filtering
Pulse compression · SCNR gain

CLASSICAL · MTI
Delay-line canceller
Stationary-clutter rejection · Doppler high-pass

CLASSICAL · CI / MTD
Fourier Doppler filtering
Slow-time integration · range-Doppler map

CLASSICAL · CFAR
Detection & decision
Threshold / local background

Reference Init · conjugate-reversed transmit waveform
LDC Init · binomial delay taps
DF Init · Fourier cosine/sine kernels
CFAR-inspired · local peak-contrast kernel

ECHMO · 01
LMF
K parallel chains · reference-seeded complex conv

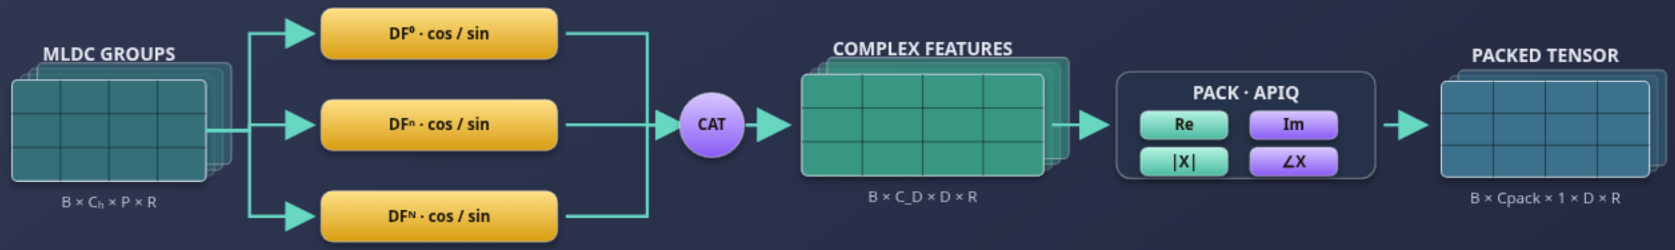
ECHMO · 02
MLDC
N+1 parallel orders · padding · delay cancellation · channel concat

ECHMO · 03
GDF
Grouped Doppler filters · cos/sin init · learnable offsets

ECHMO · 04
PKC-Seg
PeakConv encoder → decoder → dense mask

GDF · Grouped Doppler filters

Fourier-seeded cosine / sine filters are grouped and learnable offsets adapt their responses.



$$D^{\star\dagger}\big(\{\tilde{X}^n\};\{D^{n\dagger}\}\big) = \dot{X}^0 \oplus \cdots \oplus \dot{X}^N, \quad \dot{X}^n = D^{n\dagger}\big(\tilde{X}^n; D^{n\dagger}\big)$$

DF INIT

$\{\cos, \sin\}$ Fourier bases $+$ learnable frequency offsets

THE METHOD · REINTERPRETATION

From raw echoes to semantic masks: the EchMo architecture

Physical priors are no longer fixed code: they become initialization and differentiable computation flow, adapted by data within a physics-shaped architecture.

input: B×C×P×R complex I/Q tensor

▶ ANIMATE ARCHITECTURE

Forward signal flow in both pipelines — Classical operator reinterpreted as a learnable module

CLASSICAL · MF
Matched filtering
Pulse compression · SCNR gain

→

CLASSICAL · MTI
Delay-line canceller
Stationary-clutter rejection · Doppler high-pass

→

CLASSICAL · CI / MTD
Fourier Doppler filtering
Slow-time integration · range-Doppler map

→

CLASSICAL · CFAR
Detection & decision
Threshold / local background

Reference Init · conjugate-reversed transmit waveform

LDC Init · binomial delay taps

DF Init · Fourier cosine/sine kernels

CFAR-inspired · local peak-contrast kernel

ECHMO · 01
LMF
K parallel chains · reference-seeded complex conv

→

ECHMO · 02
MLDC
N+1 parallel orders · padding · delay cancellation · channel concat

→

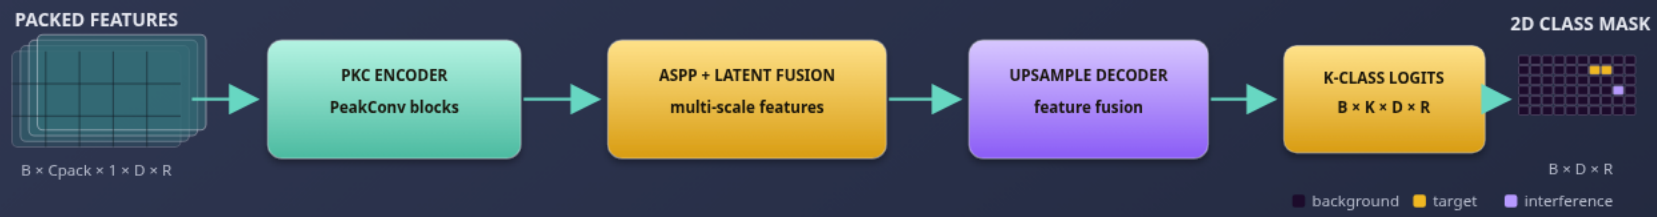
ECHMO · 03
GDF
Grouped Doppler filters · cos/sin init · learnable offsets

→

ECHMO · 04
PKC-Seg
PeakConv encoder → decoder → dense mask

PKC-Seg · PeakConv segmentation head

PKCIn4ANY encoder → ASPP / latent fusion → upsampling decoder → class logits.



$$Z = \mathrm{Pack}\big(D^{*\dagger}(\tilde{X})\big), \quad Y = \mathrm{PKCIn4ANY}_{cuda}(Z) \in \mathbb{R}^{B\times K\times D\times R}$$

TENSOR PATH

$$B\times C_{pack}\times 1\times D\times R \;\rightarrow\; B\times K\times D\times R \text{ logits} \;\rightarrow\; B\times D\times R \text{ class mask}$$

RADAR ECHO DATASET

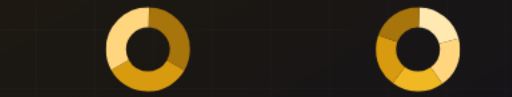
Two complementary raw-echo datasets, from controlled simulation to hardware-in-the-loop.

SYN-RED · EM SIMULATION · LOW-SLOW-SMALL TARGETS



1,000 frames per (pattern × range × SNR) · **27,000** frames / 27 sequences total

HIL-RED · HARDWARE-IN-THE-LOOP · ACTIVE JAMMING



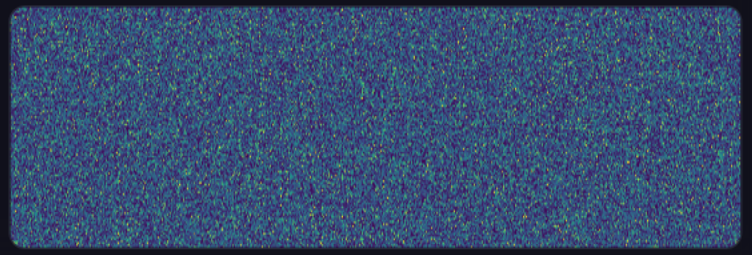
10,000 frames per (jamming type × JSR) · **150,000** CPIs / 15 sequences total

Dataset specification	Syn-RED · electromagnetic simulation	HiL-RED · hardware-in-the-loop
Task & construction	Low-slow-small target detection; signal-level EM simulation	Target–jammer recognition; controller and IF acquisition loop
Band / waveform	X-band · 9.8 GHz · LFM 16 μs / 8 MHz	Pulse 4 μs / 40 MHz
Pulse structure	PRI 100 μs × 128 pulses	PRI 10.24 μs × 128 pulses
Baseband sampling	10.24 MHz	100 MHz
Scenario dimensions	Line / S-turn / Orbit · 2 / 3 / 5 km · 10 / 14 / 18 dB	NBJ / WBJ / DFTJ / ISRJ / SIJ · −20 / −15 / −10 dB JSR
Paper scale	27 sequences · 27,000 frames	15 sequences · 150,000 CPIs; HiL-RED30K is a uniform 20% subset

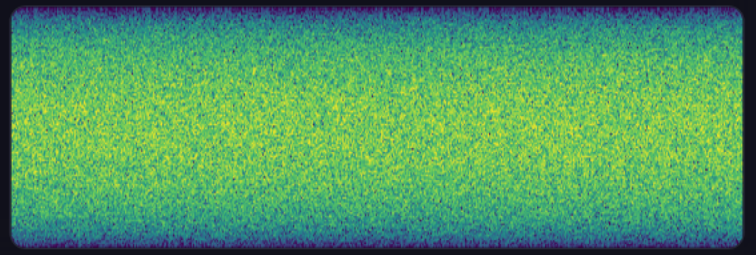
DATASET **Syn-RED · simulated** | HiL-RED · hardware-in-the-loop | RANGE 5000 m | TARGET S-shape | SNR +10 dB

Samples on this page are read directly from the project-local datasets/Syn-RED and datasets/HiL-RED.

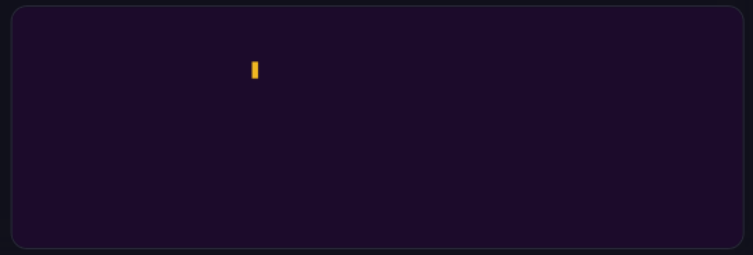
RAW I/Q ECHO



MTD / RANGE–DOPPLER



DENSE ANNOTATION · ORIGINAL LABEL IDs



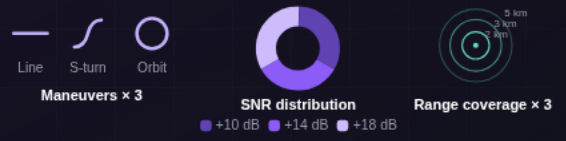
Syn-RED · 5000 m · S-shape · +10 dB · frame 000000

datasets/Syn-RED/data_5000m_S_10dB · source frame 000000 · ID 0 · bg · 99.98% · ID 1 · 0.02%

RADAR ECHO DATASET

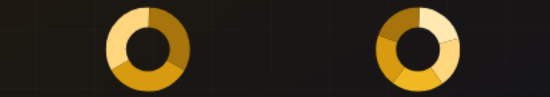
Two complementary raw-echo datasets, from controlled simulation to hardware-in-the-loop.

SYN-RED · EM SIMULATION · LOW-SLOW-SMALL TARGETS



1,000 frames per (pattern × range × SNR) · **27,000** frames / 27 sequences total

HIL-RED · HARDWARE-IN-THE-LOOP · ACTIVE JAMMING



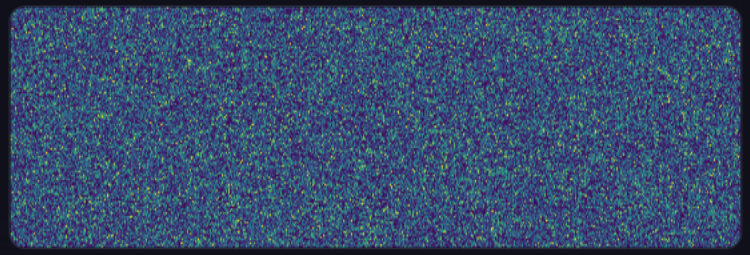
10,000 frames per (jamming type × JSR) · **150,000** CPIs / 15 sequences total

Dataset specification	Syn-RED · electromagnetic simulation	HiL-RED · hardware-in-the-loop
Task & construction	Low-slow-small target detection; signal-level EM simulation	Target–jammer recognition; controller and IF acquisition loop
Band / waveform	X-band · 9.8 GHz · LFM 16 μs / 8 MHz	Pulse 4 μs / 40 MHz
Pulse structure	PRI 100 μs × 128 pulses	PRI 10.24 μs × 128 pulses
Baseband sampling	10.24 MHz	100 MHz
Scenario dimensions	Line / S-turn / Orbit · 2 / 3 / 5 km · 10 / 14 / 18 dB	NBJ / WBJ / DFTJ / ISRJ / SIJ · −20 / −15 / −10 dB JSR
Paper scale	27 sequences · 27,000 frames	15 sequences · 150,000 CPIs; HiL-RED30K is a uniform 20% subset

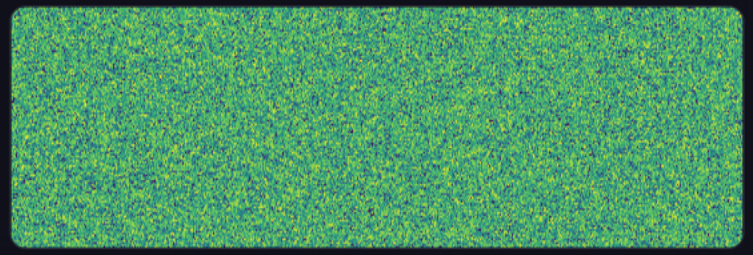
DATASET | Syn-RED · simulated | **HiL-RED · hardware-in-the-loop** | JAMMING TYPE | NBJ · Narrow-band barrage | JSR | **-20 dB**

Samples on this page are read directly from the project-local datasets/Syn-RED and datasets/HiL-RED.

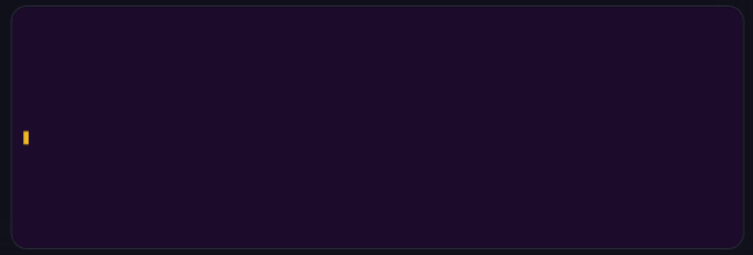
RAW I/Q ECHO



MTD / RANGE–DOPPLER



DENSE ANNOTATION · ORIGINAL LABEL IDs



HiL-RED · NBJ · Narrow-band barrage · -20 dB · frame 202509091239203879

datasets/HiL-RED/barrage_narrow_jamming_-20dB · source frame 202509091239203879 ID 1 · 0.01% ID 6 · bg · 99.99%

RECK · RELAXING PRIOR CONSTRAINTS

ReCK: how much prior should remain?

LMF · ReferenceInit
kernel ← conjugate-flipped reference waveform — at init, LMF is exactly the classical matched filter

MLDC · LDCInit
taps ← binomial solution $h_i = (-1)^i \cdot C(n, i)$ — N+1 parallel orders, identity path included

GDF · DFInit
filters ← Fourier cos / sin bases (Euler) — training frees the uniform Doppler bins into non-uniform offsets

▶ PLAY

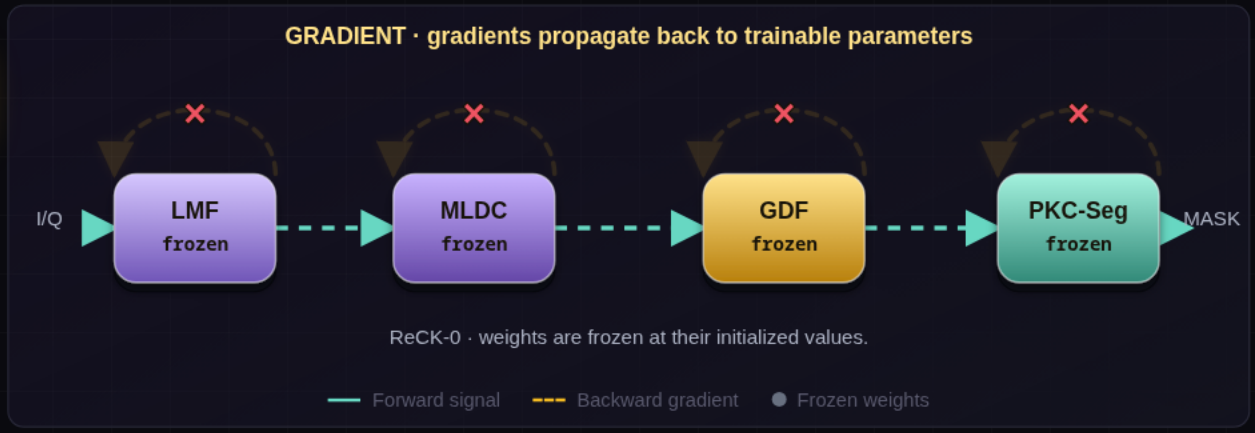
ReCK-0 · frozen
initialized, never updated

ReCK-1 · cautious FT
$\eta^* = \alpha\eta$ — small steps away

ReCK-2 · full FT (default)
joint training, global η

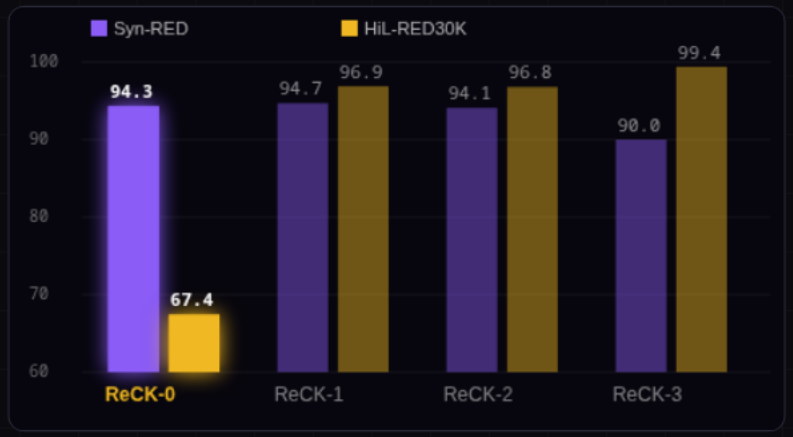
ReCK-3 · structure only
compute flow kept, init discarded

ReCK GRADIENT FLOW · FORWARD SIGNAL & BACKWARD UPDATE



Teal marks forward signals; yellow curves return from each module's output to its own input. In module-wise mode, the selected module is ReCK-0 and all others are ReCK-2.

Global | Module-wise



frozen prior: fine in simulation (94.3), collapses in reality (67.4).

More prior is not always better: waveform correlation makes matched filtering the most robust prior; uniform Doppler bins fail first under jamming. ReCK's practical lesson: the data domain determines how much prior to retain.

RECK · RELAXING PRIOR CONSTRAINTS

ReCK: how much prior should remain?

LMF · ReferenceInit
kernel ← conjugate-flipped reference waveform — at init, LMF is exactly the classical matched filter

MLDC · LDCInit
taps ← binomial solution $h_i = (-1)^i \cdot C(n, i)$ — N+1 parallel orders, identity path included

GDF · DFInit
filters ← Fourier cos / sin bases (Euler) — training frees the uniform Doppler bins into non-uniform offsets

▶ PLAY

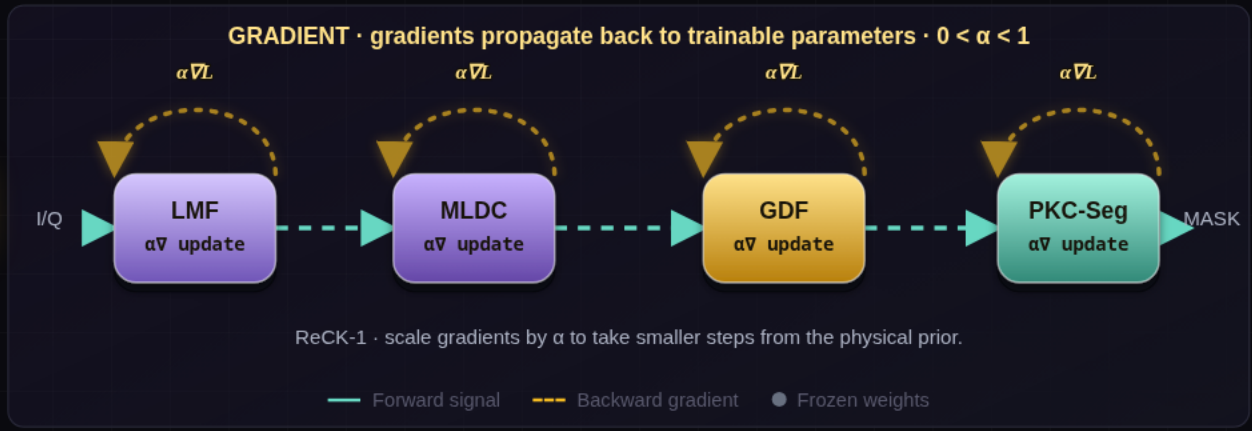
ReCK-0 · frozen
initialized, never updated

ReCK-1 · cautious FT
η* = αη — small steps away

ReCK-2 · full FT (default)
joint training, global η

ReCK-3 · structure only
compute flow kept, init discarded

ReCK GRADIENT FLOW · FORWARD SIGNAL & BACKWARD UPDATE



Teal marks forward signals; yellow curves return from each module's output to its own input. In module-wise mode, the selected module is ReCK-0 and all others are ReCK-2.

Global | Module-wise



cautious FT: best in simulation (94.7) — the more idealized the data, the more the prior deserves trust.

More prior is not always better: waveform correlation makes matched filtering the most robust prior; uniform Doppler bins fail first under jamming. ReCK's practical lesson: the data domain determines how much prior to retain.

RECK · RELAXING PRIOR CONSTRAINTS

ReCK: how much prior should remain?

LMF · ReferenceInit
kernel ← conjugate-flipped reference waveform — at init, LMF is exactly the classical matched filter

MLDC · LDCInit
taps ← binomial solution $h_i = (-1)^i \cdot C(n, i)$ — N+1 parallel orders, identity path included

GDF · DFInit
filters ← Fourier cos / sin bases (Euler) — training frees the uniform Doppler bins into non-uniform offsets

▶ PLAY

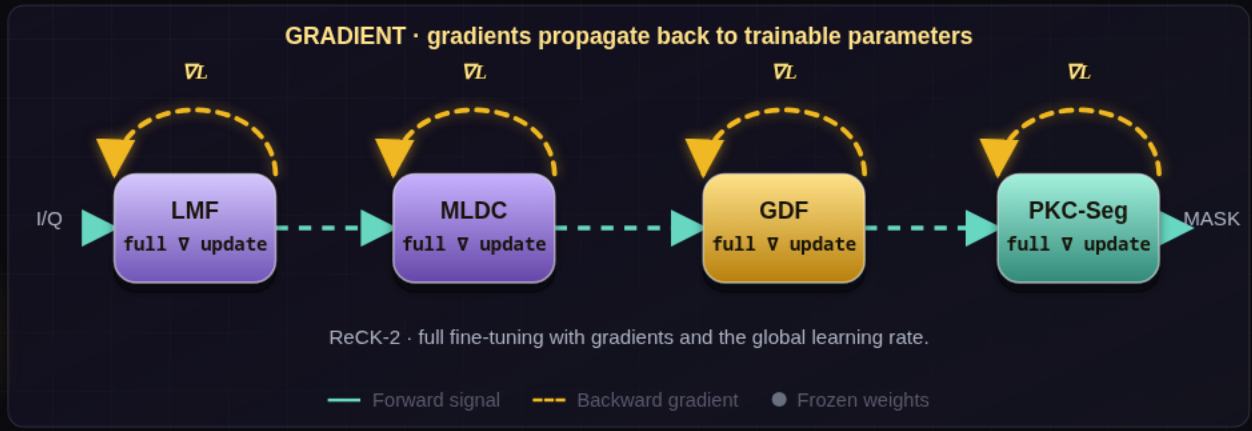
ReCK-0 · frozen
initialized, never updated

ReCK-1 · cautious FT
η* = αη — small steps away

ReCK-2 · full FT (default)
joint training, global η

ReCK-3 · structure only
compute flow kept, init discarded

ReCK GRADIENT FLOW · FORWARD SIGNAL & BACKWARD UPDATE



Teal marks forward signals; yellow curves return from each module's output to its own input. In module-wise mode, the selected module is ReCK-0 and all others are ReCK-2.

Global | Module-wise



full FT (default): the balanced compromise (94.1 / 96.8).

More prior is not always better: waveform correlation makes matched filtering the most robust prior; uniform Doppler bins fail first under jamming. ReCK's practical lesson: the data domain determines how much prior to retain.

RECK · RELAXING PRIOR CONSTRAINTS

ReCK: how much prior should remain?

LMF · ReferenceInit
kernel ← conjugate-flipped reference waveform — at init, LMF is exactly the classical matched filter

MLDC · LDCInit
taps ← binomial solution $h_i = (-1)^i \cdot C(n, i)$ — N+1 parallel orders, identity path included

GDF · DFInit
filters ← Fourier cos / sin bases (Euler) — training frees the uniform Doppler bins into non-uniform offsets

▶ PLAY

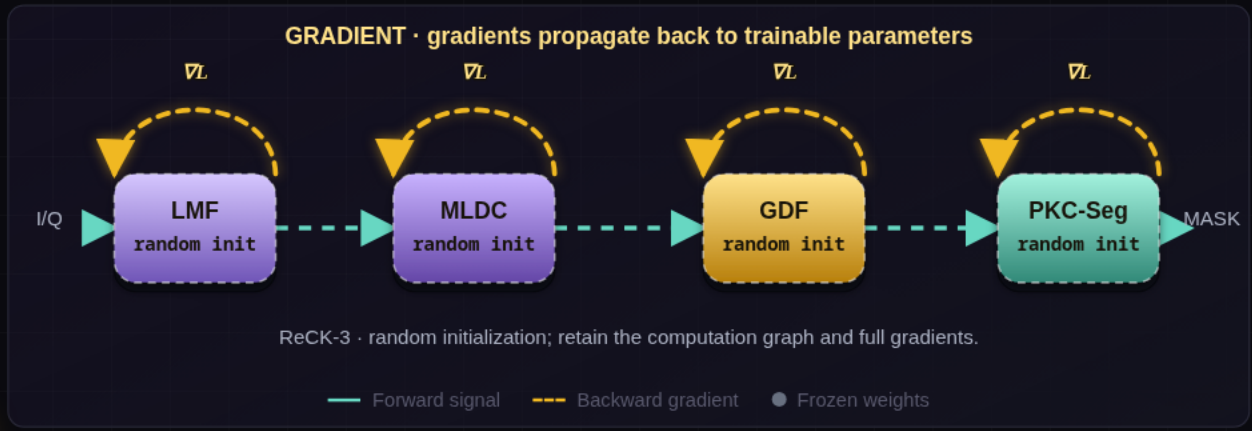
ReCK-0 · frozen
initialized, never updated

ReCK-1 · cautious FT
η* = αη — small steps away

ReCK-2 · full FT (default)
joint training, global η

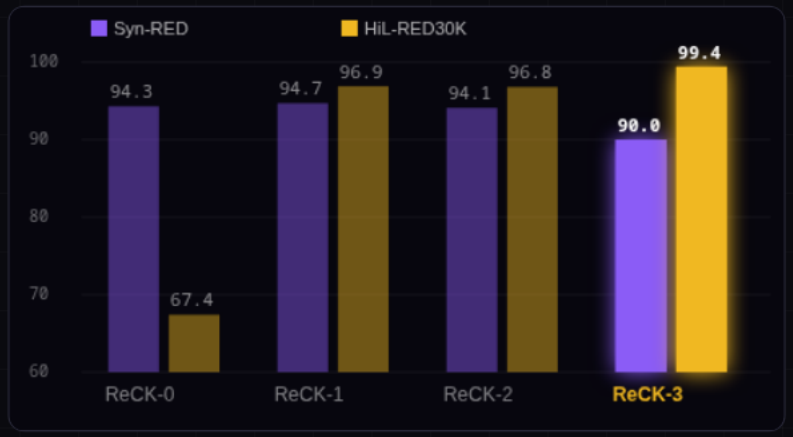
ReCK-3 · structure only
compute flow kept, init discarded

ReCK GRADIENT FLOW · FORWARD SIGNAL & BACKWARD UPDATE



Teal marks forward signals; yellow curves return from each module's output to its own input. In module-wise mode, the selected module is ReCK-0 and all others are ReCK-2.

Global | Module-wise



structure only (random init): best under real jamming (99.4), weakest in simulation (90.0).

More prior is not always better: waveform correlation makes matched filtering the most robust prior; uniform Doppler bins fail first under jamming. ReCK's practical lesson: the data domain determines how much prior to retain.

RECK · RELAXING PRIOR CONSTRAINTS

ReCK: how much prior should remain?

LMF · ReferenceInit
kernel ← conjugate-flipped reference waveform — at init, LMF is exactly the classical matched filter

MLDC · LDCInit
taps ← binomial solution $h_i = (-1)^i \cdot C(n, i)$ — N+1 parallel orders, identity path included

GDF · DFInit
filters ← Fourier cos / sin bases (Euler) — training frees the uniform Doppler bins into non-uniform offsets

▶ PLAY

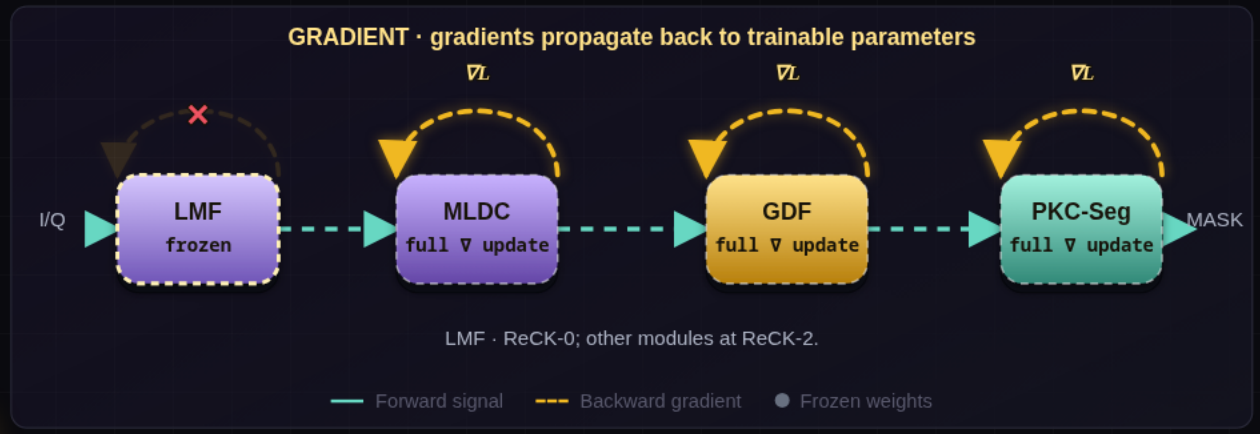
ReCK-0 · frozen
initialized, never updated

ReCK-1 · cautious FT
η* = αη — small steps away

ReCK-2 · full FT (default)
joint training, global η

ReCK-3 · structure only
compute flow kept, init discarded

ReCK GRADIENT FLOW · FORWARD SIGNAL & BACKWARD UPDATE



Teal marks forward signals; yellow curves return from each module's output to its own input. In module-wise mode, the selected module is ReCK-0 and all others are ReCK-2.

Global | Module-wise

LMF | MLDC | GDF



LMF at ReCK-0; MLDC / GDF at ReCK-2.

More prior is not always better: waveform correlation makes matched filtering the most robust prior; uniform Doppler bins fail first under jamming. ReCK's practical lesson: the data domain determines how much prior to retain.

EXPERIMENTS · MAIN RESULTS

Two settings, distinct challenges, one decisive set of results

1.29M
parameters — 1% of FCN-8s

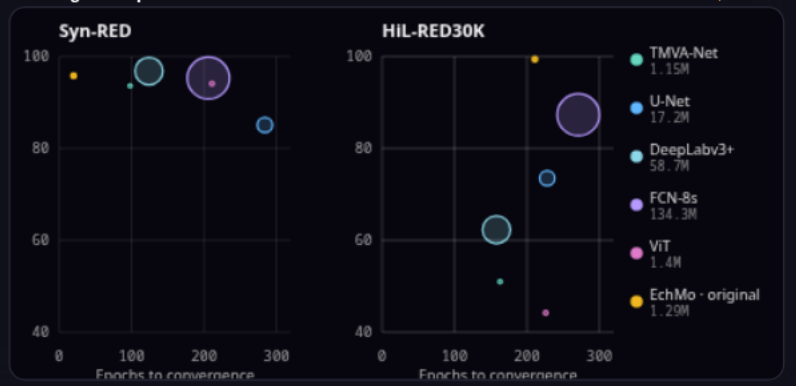
97.6 / 98.8
avg mIoU / mDice across both benchmarks

0.999
ROC AUC across three data settings (Fig. 28)

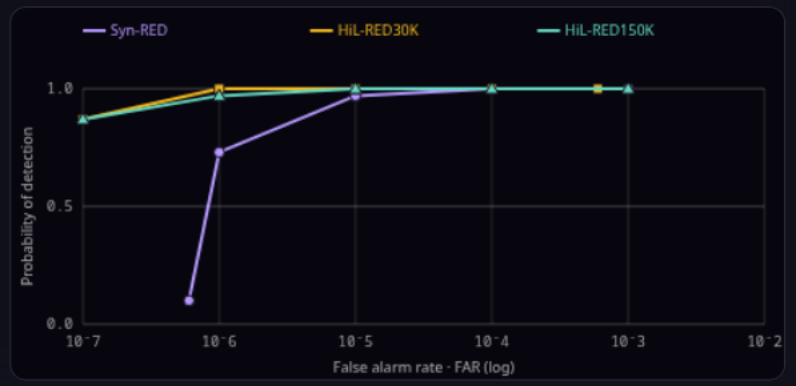
Architecture	Params	Syn-RED mIoU / mDice	HiL-RED30K mIoU / mDice
TMVA-Net	1.15M	93.6 / 96.6	51.0 / 64.5
U-Net	17.2M	85.1 / 91.2	73.5 / 81.9
DeepLabv3+	58.7M	96.8 / 98.3	62.3 / 71.9
FCN-8s	134.3M	95.3 / 97.6	87.3 / 92.1
ViT	1.40M	94.1 / 96.9	44.2 / 56.2
EchMo · unaccelerated	1.29M	95.8 / 97.8	99.4 / 99.7

mIoU / mDice (%) · baselines use RSP intermediate inputs; EchMo uses raw I/Q

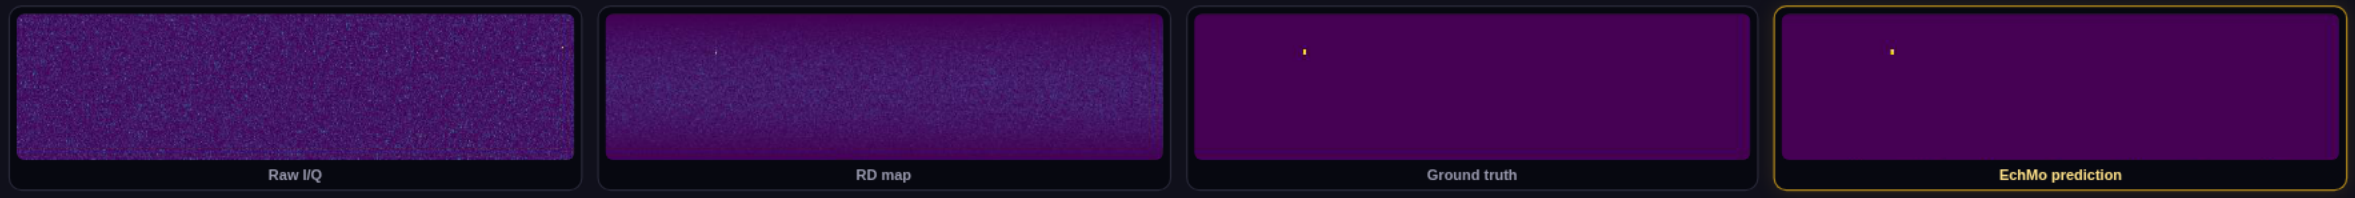
Convergence epochs × mIoU



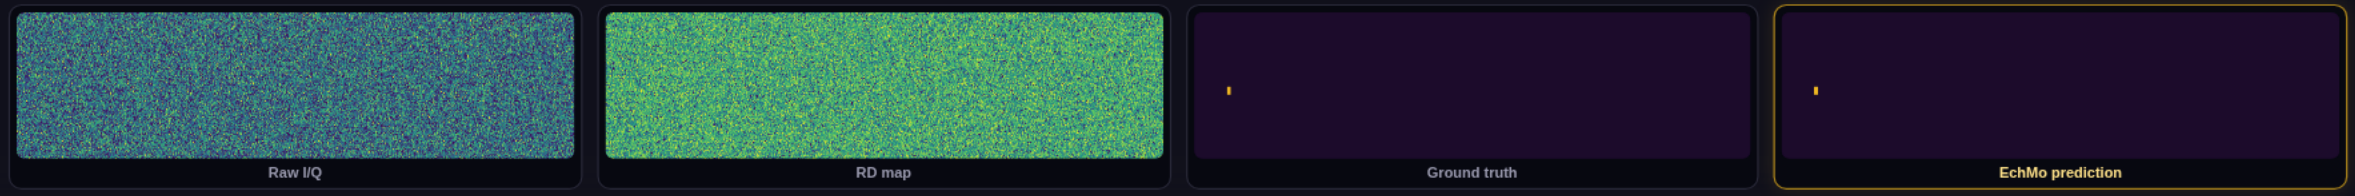
Detection ROC



Syn-RED · EchMo semantic segmentation example



HiL-RED · wide-band barrage jamming (WBJ) · −20 dB JSR



EXPERIMENTS · ACCELERATION & SCALING

Optimize the compute graph, then scale in three directions.

LMF · V7S

4 real convs → 1 packed conv

Re×Re | Im×Im | Re×Im | Im×Re → packed complex conv

Preserves the complex mapping; retains serial chains to avoid V7P's memory and numerical trade-offs.

MLDC · PAR

2(N+1) launches → 2

order 0 | order 1 | … order N → kernel bank

Pad order kernels to a common height, pack, and execute all orders in parallel.

GDF · EINSUM

Grouped projection by einsum

P pulses | DF cos/sin → grouped projection

Projects directly over pulses; forward and input gradients match, with lower peak memory.

ECHMO → ECHMO†

5.083× **1.266×**

Forward 1041.676 → 204.889 ms · backward 746.699 → 589.958 ms

−62.38%

Backward peak memory 1394.54 → 524.61 MB

SCALING ARCHITECTURES

same EchMo† backbone · three complementary extensions

01 · DEPTH **EchMoDeep**



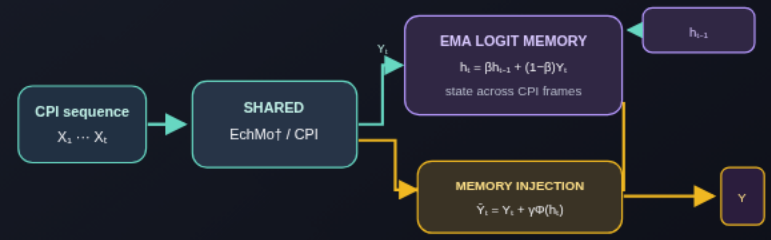
Two parallel EchMo† backbones operate at full and reduced resolution; upsampled logits are added and fused with a 1×1 convolution.

02 · WIDTH **EchMo4MoE**



Four complete EchMo† experts process each input; a statistics gate estimates sample weights for top-2 logit fusion (all experts are computed).

03 · TEMPORAL **EchMoMemo**



Reuse EchMo† per CPI; update EMA state from current logits and inject its correction without changing the physics frontend.

HiL-RED150K · SCALING RESULTS

vertical comparison of the EchMo† backbone and three extensions

Model	Params	GFLOPs	Latency ms/iter	Memory MB/GPU	mIoU	mDice	Epoch
EchMo† backbone	~9.3M	~81.2	1325.9	21978.6	93.4%	96.3%	48/50
EchMoDeep depth	~16.0M	~121.8	2410.6	37222.2	99.5%	99.7%	50/50
EchMo4MoE width	~32.0M	~324.8	4288.6	71175.8	98.6%	99.3%	41/50
EchMoMemo temporal	~9.3M	~81.2	2494.0	42169.6	97.1%	98.5%	50/50

Scaling runs use the EchMo† backbone and amplitude feature packing; platforms / batch sizes differ, so latency is contextual, not a direct cross-model comparison. Values from the paper appendix.

TAKEAWAY

Echo in, mask out.

① EchMo achieves end-to-end radar semantic segmentation from raw I/Q echoes.

Pulse compression, moving-target cancellation, and coherent integration are rebuilt as the differentiable, learnable LMF–MLDC–GDF–PKC-Seg modules, embedding radar physics into a deep network through computation flow and parameter initialization.

② The ReCK ablation framework.

Quantifies how prior-constraint strength shapes the model: simulated scenes favor strong priors, while complex electromagnetic scenes with active jamming call for relaxed priors and data-driven adaptive learning.

③ Equivalent compute-graph substitution and three-axis scaling.

Over 5× faster forward inference with markedly lower memory; EchMoDeep / EchMo4MoE / EchMoMemo extend the backbone along the orthogonal depth, width, and temporal axes, demonstrating scalable capacity.

④ Syn-RED and HiL-RED: two radar I/Q echo datasets.

Targeting low-slow-small simulation and hardware-in-the-loop active jamming respectively, they support echo-level radar semantic perception; baseline comparisons confirm EchMo's segmentation advantage under complex jamming.

A small step for radar perception — a giant step for remote perception: bringing intelligence to farther physical space.

CODE · DATASETS · DEMO
github.com/zlw9161/PeCoRSS
Syn-RED / HiL-RED datasets and all visualizations

CHECKPOINTS · VISUALIZATIONS
huggingface.co/PeCoLab/EchMo-Vis
Model weights: PeCoLab/EchMo-ckpts

See you at the poster!