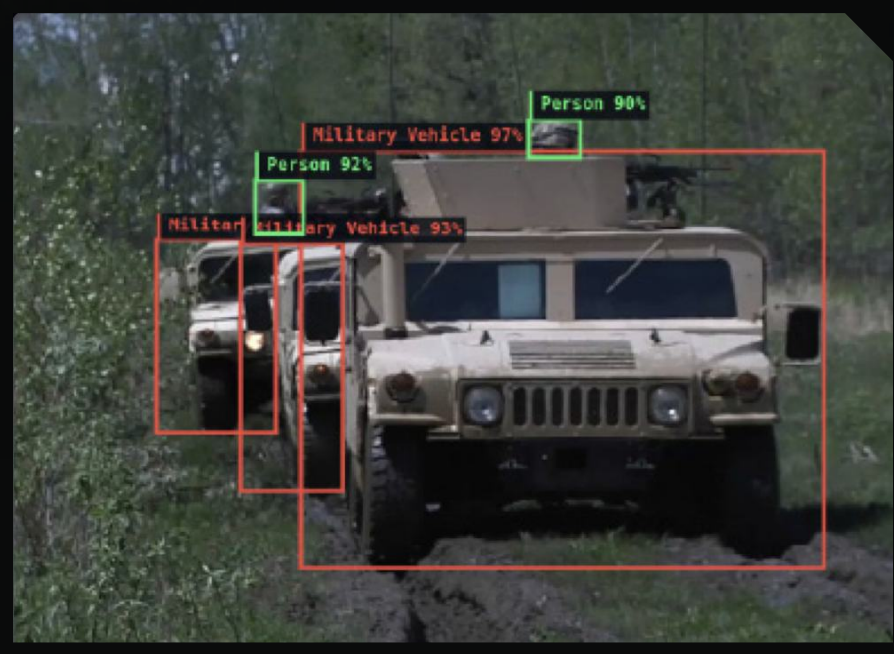
PhantomSIGHT™

From Raw Footage To Fielded ATR

What Tailoring Buys You

A mission-tailored model against a general-purpose one. On a frozen, out-of-distribution set of 1,160 images and 1,460 armored vehicles, drawn from sources absent from either model's training, PhantomSIGHT™'s 3.0-million-parameter model was evaluated against YOLOv8x-World v2, a 72.9-million-parameter open-vocabulary detector built to find anything.

2.5×	**Detection Rate** The detection rate of the general-purpose detector, at the threshold operators use.
~½	**Military vs Civilian Confusion** Cut nearly in half: tailored classes separate what the mission cares about.
24.2× fewer	**Parameters** 3.0 million against 72.9 million, small enough to run at the edge.



▲ PhantomSIGHT™ run-model output on operator-supplied video: military vehicles and personnel detected and classified frame by frame.

Deployed at the Edge

Validated on ModalAI VOXL and NVIDIA Jetson ORIN compute.
Integrates through ONIX, Tensor RT, and ROS2. Full pipeline runs air-gapped.

25 FPS	80.3%	23 MB	3.0M
Real-time onboard inference	Detection accuracy	Deployed model footprint	Model parameters

© VICTUS Technologies, Inc. (CAGE: 9UH37)

VICTUS

What PhantomSIGHT™ is

One platform that takes an operator from raw footage to a fielded Automatic Target Recognition (ATR) model. Annotate data, tailor the ATR to your mission set and domain, train and test models, and reuse, adapt, and modify datasets, all in one place, running on Commercial Off-The-Shelf (COTS) edge compute. No specialized machine-learning team, no cloud dependency, no model black box.

Why the Platform Matters

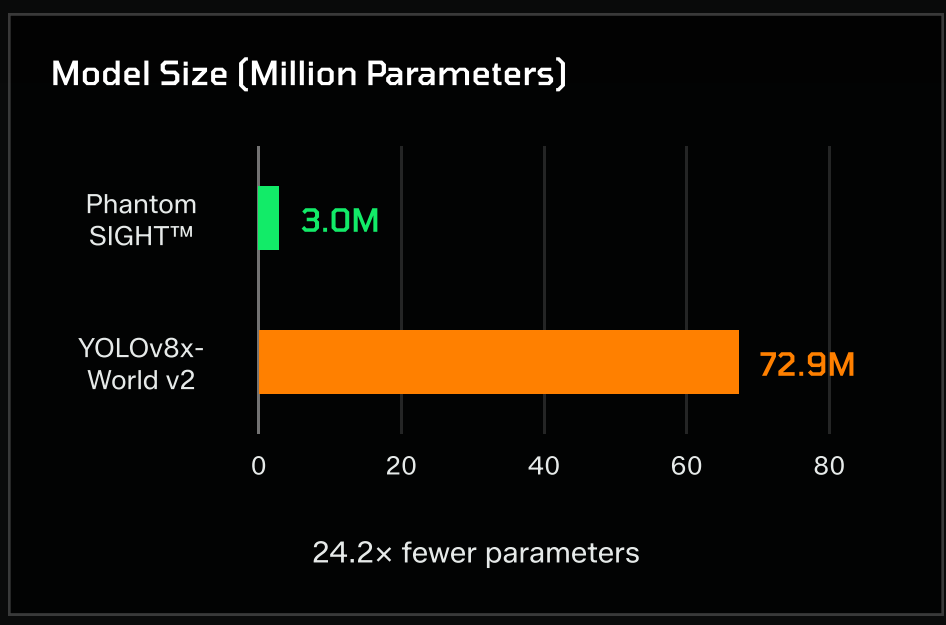
A detector alone is not a capability. What makes ATR useful in the field is how quickly it can be retargeted to a new target set, theater, or sensor, and kept current as the mission changes. PhantomSIGHT™ puts that loop in the operator's hands, with low cost cameras, Autonomous Aerial Systems (AAS), ground robotics, and surface and sub-surface vessels.

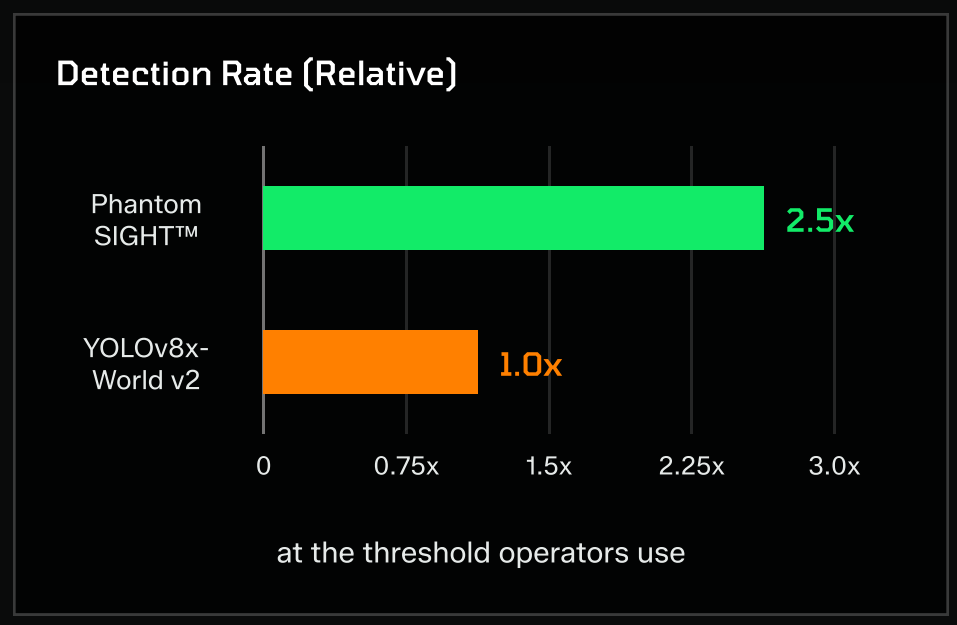
One Platform, All in One Place

01	02	03	04	05
Annotate Data	**Build Datasets**	**Tailor to the Mission**	**Train and Test Models**	**Reuse, Adapt, Modify**
Label in the field and let labels carry across frames.	Pull sensor feeds and open-source footage into one dataset.	Define target and background classes for your mission and domain.	Fine-tune from as few as five images and validate before fielding.	Remap, merge, or drop classes to reuse datasets.

Size vs. Performance

PhantomSIGHT™, mission-tailored, 3.0M parameters
YOLOv8x-World v2, general-purpose open-vocabulary detector, 72.9M parameters





Against 24.2 times the parameter count, PhantomSIGHT™ achieved 2.5 times the detection rate at the threshold operators use and cut military-versus-civilian confusion nearly in half. Both models were scored on the same frozen set of 1,160 images, none of it drawn from either model's training sources. Detection rate is shown relative to YOLOv8x-World v2 (1.0×).

NEXT STEP: SCHEDULE A CAPABILITY BRIEF

getvictus.si | info@getvictus.ai

© VICTUS Technologies, Inc. (CAGE: 9UH37)