BRIEFING B3 UNCLASSIFIED

Forest-to-Desert Framework

Topic development for ARL-BAA-0107 using ARL-BAA-0038: the framework defined, its seven-phase protocol, and 60 research topics across six Army technology categories

John Kolb
SOL Engineering Services

PUBLISHED
November 2026

READING TIME
26 min

Executive Summary

This report develops candidate research topics for ARL-BAA-0107, *Academic Applied and Advanced Innovation*, using the research goals of ARL-BAA-0038, *Artificial Intelligence and Machine Learning with Extremely Sparse Data*, as the technical foundation. Elements of ARL-BAA-0064, *Autonomous Sensing and Information Fusion for Advanced Indications and Warnings*, are also applied. SOL Engineering Services and its university partner bring one architecture to every topic: the Forest-to-Desert Framework™, built on SOL Aquifer™. The report has three parts (Table 1).

Table 1. Report overview

Part	What it covers	Main takeaway
A. AI/ML research direction (ARL-BAA-0038)	ARL's five research thrusts; key concepts; the Forest-to-Desert Framework™: SOL Aquifer, uncertainty-aware neuro-symbolic inference and the data flywheel, carried out through a seven-phase protocol	Codified physics and expert rules reduce the need for data; separating epistemic and aleatoric uncertainty decides when to trust learned models; a data flywheel improves both as real data arrives
B. Topic portfolio	60 topic ideas across six Army technology categories, each paired with a university role and a product	32 topics are strong fits for SOL's existing AI work
C. SOL Aquifer™ fit across six categories	SOL Aquifer™ mapped onto every category; university mechatronics and additive manufacturing lab use	Most topics fit where a slow physics solver exists; the university lab supplies the real calibration data

3	7	60
PARTS OF THE FRAMEWORK: SOL AQUIFER, INFERENCE, DATA FLYWHEEL	PHASES IN THE FOREST-TO-DESERT PROTOCOL	TOPIC IDEAS ACROSS SIX ARMY TECHNOLOGY CATEGORIES

Part A. AI/ML Research Direction

1. ARL-BAA-0038 research goals

ARL-BAA-0038 seeks research that lets Army platforms, intelligence systems and command systems learn, infer and predict under extreme epistemic uncertainty. It expects data-driven learning combined with domain-expert rules to reduce the need for data, and uncertainty-aware neuro-symbolic processing to expose the limits of what data-driven models can infer. The research outline developed for this topic has five ARL research thrusts (T1–T5), which are used throughout this report (Table 2).

Table 2. ARL research thrusts

Thrust	Focus
T1. Optimal learning with limited supervision	Expert rules as priors or constraints; separating aleatoric and epistemic uncertainty; sample-complexity bounds; handling imperfect or conflicting expertise
T2. Efficient approximations	Variational inference, ensembles, evidential learning; knowledge compilation for the symbolic side; bounding the gap to optimal
T3. When traditional AI/ML is enough	Detecting regime changes; switching between data-driven, hybrid and rule-dominant modes; guarantees across switches
T4. Edge processing	Compression that preserves calibrated uncertainty; SWaP-aware and anytime inference; edge–cloud partitioning in DDIL conditions
T5. Distributed learning with GNNs	Uncertainty-carrying message passing; federated learning under intermittent comms; trust-aware fusion

Cross-cutting elements are an expert elicitation method; Army-relevant testbeds (ISR sensor fusion, C2 decision support, autonomy in unfamiliar terrain); and metrics covering accuracy versus labeled samples, calibration, out-of-distribution detection, and latency and energy on target hardware.

2. Key concepts

- **Epistemic uncertainty** comes from lack of knowledge, such as a model that has never seen “desert” terrain. More data or expert rules reduce it.
- **Aleatoric uncertainty** is randomness built into the problem, such as sensor noise or a blurry image at long range. More data does not remove it.
- **Novelty** is how far a new input sits from the conditions already covered. It is a warning sign of epistemic uncertainty, not the uncertainty itself.
- **Error** is the prediction minus the measured value. It is known only after measurement, and it is used to check that the stated uncertainty was honest.
- **An AI agent** (an autonomous software system) pursues a goal over several steps: it perceives, decides, acts and observes the result within guardrails, planning and executing a workflow until it is complete. Knowing when it does not know is what lets it defer safely. Not every use of the framework is an agent; often the output is a trusted prediction or recommendation for a person to act on at critical junctures.

3. The Forest-to-Desert Framework™

The Forest-to-Desert Framework is a protocol developed and branded by SOL within a MOSA (Modular Open Systems Approach) framework for software development. It is a flexible concept that provides AI support across the landscape: design, modeling and simulation, human resources, costing, process development, logistics and more. It meets ARL’s research goals for ARL-BAA-0038 and for ARL-BAA-0064, a topic on which SOL has experience with ERDC.

What the name means. The name describes the core problem in most dynamic, data-centered solutions: capability learned where data is plentiful (the forest) must work where data is sparse (the desert), whether the desert is new terrain in the field or a new operating condition in a process.

The framework has three parts, carried out through a seven-phase protocol. Figures 1 and 2 together define the full architecture: Figure 1 shows the loop, and Figure 2 shows the seven-phase protocol. The three parts are:

01 **SOL Aquifer™ builds the knowledge.** Physics, code limits, requirements and expert rules are captured as executable software, swept across the design space and turned into surrogates that report their uncertainty. Every change is a versioned release.

02 **Uncertainty-aware neuro-symbolic inference makes every prediction.** Codified physics and learned surrogates are combined on each query, and uncertainty decides which to trust. This level and the data flywheel both pass through the SOL Aquifer foundation into a single database with provenance.

03 **The data flywheel improves both.** It monitors continuously, learns in batches and changes the system only through SOL Aquifer.

Figure 1. The Forest-to-Desert Framework™: one loop from knowledge to prediction to improvement

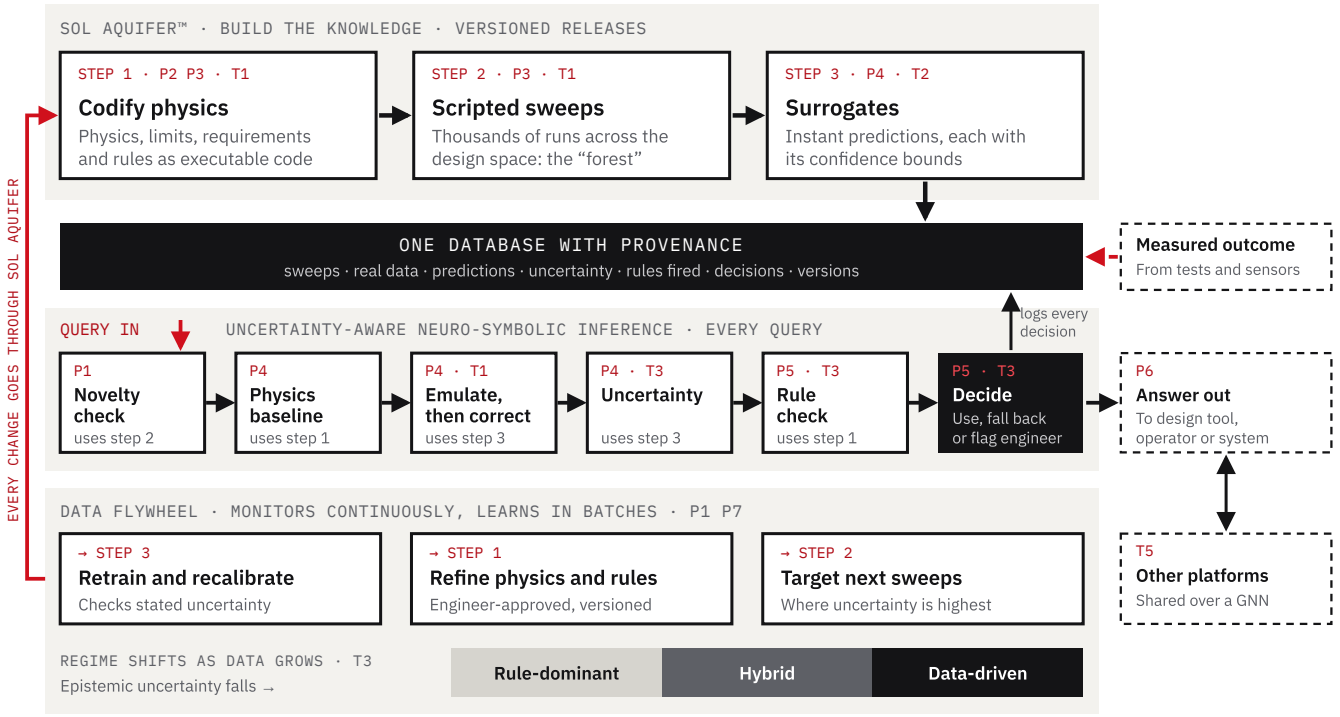

Tags show protocol phases (P1–P7), the SOL Aquifer step each run-time step uses (steps 1–3) and ARL thrusts (T1–T5). Redrawn from the report figure.

Read Figure 1 from the top. SOL Aquifer turns physics and rules into code (step 1), sweeps that code across the design space to produce thousands of labeled runs (step 2) and trains surrogates that answer instantly with a stated confidence (step 3). Everything it produces lands in one database with provenance. In the middle band, each new query runs through six steps and ends in a decision, which sends an answer out and is logged. Measured outcomes from tests and sensors arrive later and enter the same database. In the bottom band, the data flywheel reads that record and sends every improvement back to SOL Aquifer as a new versioned release, never as an untracked change.

3.1 SOL Aquifer™

SOL Aquifer, formerly called “Code as Architecture”, grew out of SOL’s *Modeling HVAC System Performance: Code as Architecture Approach to AI* work, carried out with Alabama A&M University students through Georgia Tech funded internships (PIN). It codifies each engineering step as software. An engineering lane (design, conditions, physics analysis, results) becomes a software lane: a script sweeps conditions, a physics solver runs in code, results populate a SQL database with provenance, a training loop fits support vector regression and neural network surrogates, and the surrogate delivers instant predictions. In the HVAC case, Manual J/N load calculations became real-time zone load forecasts with a carbon impact estimate. The Software Pipeline of SOL Aquifer briefing walks through that run. The name reflects its role: a reservoir of codified knowledge the system can draw on when data is needed.

Four additions adapt the pipeline to ARL’s sparse-data, high-uncertainty setting, because a surrogate trained only on simulation inherits the simulation’s errors. Each is now part of the framework:

- **Quantify uncertainty.** Ensembles or Gaussian processes put confidence bounds on every prediction. This is the uncertainty step at run time.
- **Calibrate to test data.** Solver bias is corrected with the few real test points available. Measured outcomes enter the database, and the flywheel recalibrates.
- **Active learning.** New sweep points are chosen where the surrogate is least certain. This is the flywheel’s “target next sweeps” step.
- **Envelope guardrail.** Any query outside the swept conditions is flagged instead of extrapolated. This is the novelty check at run time.

Each run-time step draws on a specific SOL Aquifer step: the novelty check uses the sweep envelope (step 2), the physics baseline and rule check use the codified physics (step 1), and the surrogate and uncertainty steps use the surrogates (step 3). Requirements, thresholds and regime settings are codified in step 1 as versioned configuration.

3.2 Uncertainty-aware neuro-symbolic inference

Every query, meaning new inputs with no answer yet, passes through six steps: a novelty check against the sweep envelope; a physics baseline from the codified model; the surrogate; uncertainty estimates, with epistemic and aleatoric uncertainty kept separate; a rule check against code limits and physical bounds; and a decision. The symbolic parts (codified physics and rules) and the neural parts (learned surrogates) combine on every prediction, and uncertainty decides which to trust.

The surrogate plays two roles over time. At Cycle 0, trained only on sweep data, it is an emulator: a fast copy of the physics that adds speed, which matters when the physics is slow, as with CFD or finite-element models. Once measured outcomes exist, it becomes a corrector that learns the gap between physics and reality. Where the physics is fast, as in HVAC load calculations, it runs live as the baseline; where it is slow, a reduced-order model or the emulator supplies the baseline.

Every decision sends an answer out to the design tool, operator or system, and every decision is logged in the database (Table 3).

Table 3. Decision paths at run time

Decision	Answer that goes out	What is logged
Use	The surrogate's prediction, with its uncertainty	Inputs, prediction, uncertainty and rules fired, marked "used"
Fall back to physics	The codified physics answer	The same record, marked "fallback", with the reason
Flag engineer	A hold, or the physics answer marked provisional	The same record, marked "flagged"; the engineer's ruling when it returns

3.3 Data flywheel

The flywheel monitors the database continuously (calibration drift, novelty rate, flagged cases) and learns in batches, on a schedule or when monitoring triggers, taking flagged and fallback cases first. It changes the system only through SOL Aquifer: retraining produces a new surrogate version (step 3); rule fixes produce a new codified model or rule version approved by an engineer (step 1); and new sweeps, or requests for physical tests, are aimed where epistemic uncertainty is highest (step 2). Small, limited updates can be made online within the guardrails, and a serious problem gets a fast versioned configuration release, such as forcing fallback to physics, without waiting for retraining.

As real data accumulates and epistemic uncertainty falls, the system moves from rule-dominant, through hybrid, to data-driven operation; a new shift sends it back toward rule-dominant mode. Deciding when to switch addresses thrust T3. Sharing predictions and uncertainty across platforms over a graph neural network addresses T5.

3.4 The seven-phase protocol: the Forest-to-Desert agent workflow

The protocol takes a detection and classification agent trained on abundant forest data to desert conditions, where data may be sparse, mostly unlabeled, and measured in the field or in a process. Figure 2 shows the seven phases, the main activities in each and the ARL thrusts each addresses. Table 4 shows where each phase sits in the framework.

Figure 2. The Forest-to-Desert agent workflow: the seven-phase protocol with ARL thrusts T1–T5

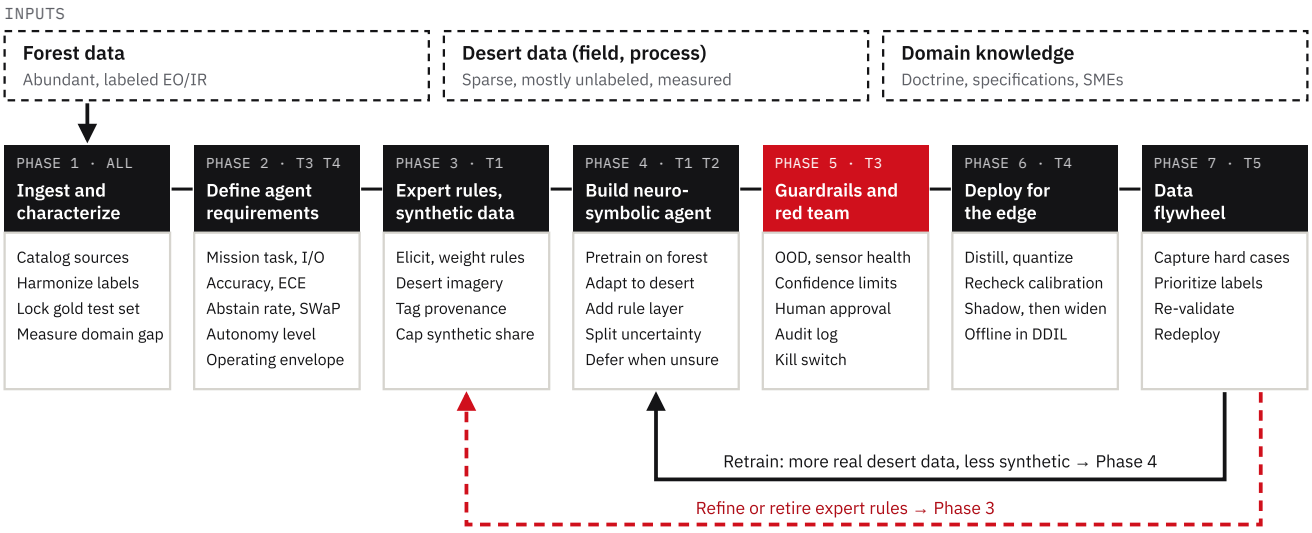

Cycle 0 relies on synthetic data and expert rules; by Cycle N, real desert data dominates and autonomy widens within the certified envelope. The AI tools used at each phase are SOL proprietary information and are not shown. Redrawn from the report figure.

Table 4. Protocol phases, activities and place in the framework

Phase	Key activities	Thrusts and place in the framework
1. Ingest and characterize	Catalog sources; harmonize labels; lock a gold desert test set; measure the domain gap	All thrusts. Novelty check for each query; flywheel monitoring across the dataset
2. Requirements	Accuracy, calibration, abstain rate, latency, SWaP; autonomy level; operating envelope	T3, T4. Codified in SOL Aquifer step 1 as versioned thresholds
3. Expert rules and synthetic data	Elicit weighted rules; simulate, style-transfer and composite desert imagery; cap synthetic weight	T1. SOL Aquifer steps 1 and 2: rules, physics and sweeps
4. Build the agent	Pretrain on forest; adapt; add rule layer; separate epistemic and aleatoric uncertainty	T1, T2. Surrogates trained in SOL Aquifer step 3 and assembled at run time
5. Guardrails	OOD detection; confidence thresholds; rule–network conflicts; human approval; audit log	T3. Rule check and decision step
6. Deploy at the edge	Distill and quantize; re-check calibration; shadow mode to wider autonomy	T4. Edge-deployable run time; answer out
7. Data flywheel	Capture uncertain and overridden cases; active learning; re-validate; redeploy	T5, with T3. Data flywheel; sharing across platforms over a graph neural network

Four design rules hold the loop together: surrogates output uncertainty, not just a number; physics stays live at run time; one database with provenance stores everything; and every change goes through SOL Aquifer. Only **three gateways** write to the database: SOL Aquifer; measured outcomes from tests, sensors and operations; and the decision step at run time.

Program cautions. CUI and classified data usually cannot reach commercial cloud APIs, so on-premises tools and open-weight models are preferred. The license of every open-source component should be checked before it is built into a deliverable. Only fielded components need verification and validation for fielding. The protocol provides gap security internally by routing everything to a single database with provenance, which would sit with SOL's AI high-performance computing center in Memphis.

Part B. Topic Portfolio Across Six Army Technology Categories

4. Program practicalities

- **Teaming.** Define SOL's and the university's roles and work shares in writing early.
- **Export control.** Hypersonics, directed energy, energetics, aviation, PNT and military communications may be ITAR or EAR controlled; settle fundamental-research status early.
- **Human subjects.** Soldier performance studies need IRB and DoW human-research review.
- **Engaging ARL.** Contact the technical points of contact named in ARL-BAA-0107 and ARL-BAA-0038, and use the white-paper step to test fit before a full proposal.
- **IP agreement.** A written allocation-of-rights agreement with the university and partners should be in place before award.

5. Topic ideas by category

Ten topics were developed for each of six technology categories (Tables 5–10). Each topic is rated by how well it fits SOL's neuro-symbolic protocol: Strong, Good or Partial, with a qualifier where a topic sits between two ratings. Part C then maps each topic onto SOL Aquifer.

5.1 Weapon Systems & Lethality

Table 5. Weapon Systems & Lethality topics

#	Topic	Research focus	University	Product	Fit
01	Hypersonic	Physics-informed ML for aerothermal loads and thermal protection response from sparse test data, with UQ	High-enthalpy facilities, CFD	Surrogate design tool with confidence bounds	Strong
02	Extreme-environment materials	Active learning to speed discovery and qualification of UHTCs and refractory alloys	Synthesis and characterization labs	Next-experiment recommender; fewer test coupons	Strong
03	Material manufacturability	In-situ additive manufacturing monitoring and ML defect prediction; process digital twin	AM machines, CT/SEM, process physics	Qualification toolkit needing fewer builds	Strong
04	Laser diagnostics	ML processing of absorption spectroscopy and PLIF data in near real time at the edge	Diagnostic rigs, spectroscopy	Embedded analysis software for test ranges	Good to Strong
05	Directed energy	Turbulence and beam-propagation models that adapt across climates	Optics labs, propagation modeling	Environment-adaptive performance predictor	Strong
06	Fire control	Uncertainty-aware sensor fusion and neuro-symbolic decision aid, human in control	Estimation, sensor fusion, human factors	Decision-support module with guardrails	Strong
07	Air & missile defense	GNN track fusion across sensor networks with degraded comms; sparse-data counter-UAS	Graph ML, networked sensing	Edge track-fusion software	Strong

#	Topic	Research focus	University	Product	Fit
08	Long-range precision fires	GPS-denied terrain and vision navigation; trajectory prediction under atmospheric uncertainty	Navigation, vision, meteorology	Edge navigation filter with uncertainty bounds	Strong
09	Energetics	ML prediction of properties, aging and stability to reduce hazardous testing	Energetics labs, computational chemistry	Formulation screening software	Good to Strong
10	Cross-cutting T&E	Bayesian design of experiments: the fewest tests that still answer the question	Statistics, uncertainty quantification	Test-planning tool for ranges and programs	Strong

5.2 Ground Vehicle Technologies

Table 6. Ground Vehicle Technologies topics

#	Topic	Research focus	University	Product	Fit
01	One-to-many C2	Workload-aware task allocation and trust calibration for one operator commanding many robotic vehicles	Human factors, human-robot interaction labs	Multi-vehicle control interface with adaptive autonomy	Partial
02	Off-road AI perception	Uncertainty-aware terrain traversability from sparse data, adapting across terrains (forest to desert)	Field robotics, computer vision	Traversability module with confidence maps	Good
03	Air-ground teaming	UAS (Drone) overhead mapping that guides UGV (Unmanned Ground Vehicle) routing through a shared world model	Multi-robot systems, SLAM (Simultaneous Location and Mapping)	Air-ground map fusion software	Good (to Strong)
04	Vehicle performance	Physics-informed digital twins of vehicle dynamics and terramechanics learned from limited test data	Terramechanics, vehicle dynamics labs	Mobility prediction tool with uncertainty bounds	Strong
05	Power management	Mission-profile energy forecasting, hybrid power control and battery health prognostics	Electrochemistry, power electronics	Energy-aware mission planner	Strong
06	Movement & maneuver	Risk-aware path planning that combines learned costmaps with doctrine-based expert rules	Motion planning, optimization	Planner with rule-based constraints	Strong
07	Long-duration ops	Self-assessment of competence: OOD detection and continual learning without forgetting over weeks	Continual learning, reliability engineering	Onboard health and competence monitor	Good

#	Topic	Research focus	University	Product	Fit
08	Human-AI collaboration	Explainable intent communication and natural-language tasking with guardrails	Cognitive science, NLP	Tasking and explanation interface	Partial
09	Human-in-the-loop ML	Learning from operator takeovers and demonstrations; active queries when the vehicle is unsure	Imitation learning, reinforcement learning	Intervention-to-training-data flywheel	Good
10	Autonomy T&E	Adversarial scenario generation and Bayesian test design for autonomy verification and validation	Formal methods, statistics	Autonomy test-scenario generator	Strong

5.3 Networking & Communications

Table 7. Networking & Communications topics

#	Topic	Research focus	University	Product	Fit
01	Resilient tactical networks	AI-driven routing and topology adaptation for mobile ad hoc networks under jamming and disconnection	Wireless networking testbeds	Adaptive routing agent for DDIL links	Strong
02	AI/ML optimization	Safe reinforcement learning for spectrum and resource allocation that stays within operator-set limits	Reinforcement learning, wireless comms	Spectrum management engine with safety constraints	Strong
03	Self-sensing networks	Degradation detection from sparse fault data that separates jamming, terrain and hardware failure, with confidence	Signal processing, network measurement	Network health monitor with confidence scores	Good
04	Self-healing networks	Neuro-symbolic fault diagnosis: operator playbooks plus learned models, with human approval for high-impact fixes	Network management, AI planning	Diagnose-and-repair agent with guardrails	Strong
05	GNN network twins	Graph neural network digital twins that predict latency and throughput before a change is applied	Graph ML, network modeling	Network digital twin for what-if testing	Strong
06	Network security	Intrusion detection that holds up under traffic shift and adversarial evasion, with few labeled attacks	Cybersecurity labs	Detector that flags novel traffic instead of guessing	Partial
07	RF spectrum awareness	Signal classification from sparse labeled RF data, augmented with synthetic signals	Software-defined radio testbeds	Edge RF classifier	Good

#	Topic	Research focus	University	Product	Fit
08	Federated learning	Communication-efficient model updates across tactical nodes on bandwidth-limited links	Distributed ML, compression	Federated learning middleware for DDIL	Good
09	Secure data sharing	Zero-trust access and lightweight, post-quantum-ready cryptography on edge devices	Cryptography, systems security	Secure edge data broker	Partial
10	Dual-use networks	Self-healing private 5G and IoT networks for disaster response, utilities and industry	5G/6G testbeds (e.g., NSF PAWR sites)	Commercial product line for Phase III	Strong

5.4 Sensing Technologies

Table 8. Sensing Technologies topics

#	Topic	Research focus	University	Product	Fit
01	PNT sensor fusion (PNT-Position, Navigation and Timing)	Uncertainty-aware fusion of inertial, vision, magnetic and other sensors that knows when its position is untrustworthy	Navigation, estimation theory	Resilient PNT filter with integrity bounds	Strong
02	Signals of opportunity	Positioning from LEO (Low Earth Orbit vs. GPS) satellite, cellular and broadcast signals when GPS is denied	Software-defined radio, satellite navigation labs	Signals-of-opportunity receiver software	Strong
03	Vision-aided nav	Matching imagery to maps across seasons and terrains (the forest-to-desert shift) with sparse reference data	Computer vision, photogrammetry	Map-matching navigation module	Good to Strong
04	Precision timing	Clock holdover with chip-scale atomic and quantum sensors, plus learned drift compensation	Atomic physics, quantum sensing labs	Timing holdover module with drift model	Good
05	GPS interference	Detect, classify and localize jamming and spoofing with ML	GNSS signal processing (Global Navigation Satellite System) vs. GPS	Spoofing and jamming detection sensor	Good
06	Collaborative PNT (Position, Navigation, Timing)	Cooperative positioning across networked units with factor graphs and GNNs over degraded comms	Multi-agent estimation	Cooperative navigation software	Strong
07	Sparse-data ATR	Target recognition that adapts to new environments using	Computer vision, remote sensing	Recognition module with	Strong

#	Topic	Research focus	University	Product	Fit
		synthetic data and expert rules		uncertainty and deferral	
08	Multi-INT fusion	Fusing EO/IR, radar, RF and acoustic data with calibrated uncertainty when sources conflict	Sensor fusion, Bayesian methods	Fusion engine with source-confidence tracking	Partial (to Good)
09	Edge ISR processing	Onboard processing that sends detections, not raw video, within SWaP and bandwidth limits	Embedded ML, computer architecture	Edge ISR processing software	Good
10	Novel sensors	Physics-aware ML for event cameras, computational imaging and SAR	Optics and radar labs	Processing pipeline for a new sensor type	Strong

5.5 Air Platforms & Aviation

Table 9. Air Platforms & Aviation topics

#	Topic	Research focus	University	Product	Fit
01	Future vertical lift	Physics-informed surrogates for rotorcraft aeromechanics and loads from limited flight test data, with UQ	Rotorcraft labs, wind tunnels, CFD	Loads surrogate tool with confidence bounds	Strong
02	Optionally manned	Autonomy that recognizes when it is outside its competence and hands control back to the pilot safely	Human factors, flight simulators	Pilot-autonomy handoff module with guardrails	Good
03	Degraded visual landing	Sensor fusion and learned perception for landing in brownout, fog and smoke	Aerial robotics, radar and lidar	Degraded-visual-environment landing aid	Strong
04	Rotorcraft health	Prognostics for rotor and drivetrain components when real failure data is sparse	Vibration analysis, reliability	Health monitoring analytics	Good
05	Swarm coordination	Decentralized coordination with GNNs that keeps working through comms loss, supervised by one operator	Multi-agent systems, graph ML	Swarm coordination software	Good
06	Airborne sensors	Onboard fusion and obstacle detection for SWaP-limited UAS payloads	Embedded sensing, computer vision	Payload processing software	Good
07	Materials	ML-assisted damage detection in lightweight composites and erosion-resistant coatings	Composites and coatings labs	Inspection analytics for maintainers	Strong
08	Components	Qualifying additively manufactured flight components with limited fatigue testing	Additive manufacturing, fatigue testing	Component qualification analytics	Strong

#	Topic	Research focus	University	Product	Fit
09	Runtime assurance	Safety monitors that bound learned flight controllers and switch to a verified backup	Formal methods, flight controls	Runtime assurance monitor	Good
10	Hybrid-electric UAS (drone)	Energy management for hybrid-electric propulsion, with dual-use for commercial drones, low weight efficient cooling system	Power electronics, propulsion labs	Energy management system for hybrid UAS	Strong

5.6 Soldier Performance

Table 10. Soldier Performance topics

#	Topic	Research focus	University	Product	Fit
01	Performance modeling	Digital twins that predict fatigue and cognitive load from sparse individual data, with uncertainty	Kinesiology, cognitive science	Performance prediction model with confidence bounds	Strong
02	LVC simulation	AI-generated scenarios, terrain and adaptive opposing-force behavior for live, virtual and constructive training	Game AI, simulation, procedural generation	Adaptive scenario generator	Strong (to Good)
03	Adaptive training	Intelligent tutoring that tracks proficiency with uncertainty and adjusts difficulty	Learning science, educational data mining	Adaptive training engine	Partial
04	C2 decision aids	Decision aids that state their uncertainty and explain recommendations, with the commander in control	Human factors, decision science	Decision-support module with guardrails	Good
05	Protective materials	ML-guided design of lightweight armor and helmets with limited impact test data	Materials, impact and blast labs	Protective design screening tool	Strong
06	Soldier power	Forecasting wearable power demand and managing batteries and energy harvesting	Electrochemistry, power electronics	Soldier power manager	Strong
07	Body-worn devices	Low-power edge AI and sensor fusion on wearables and heads-up displays	Embedded systems, HCI	Wearable edge AI software kit	Good
08	Physical performance	Injury-risk prediction from wearable sensors during load carriage and training	Biomechanics labs	Injury-risk analytics for trainers	Strong

#	Topic	Research focus	University	Product	Fit
09	Heat & fatigue	Personalized heat-strain and fatigue warnings that adapt across climates (forest to desert) from few samples	Environmental physiology	Early-warning module for wearables	Strong
10	First-responder use	Dual-use wearables and analytics for firefighters, EMS and athletes	Public health, sports science	Commercial wearable analytics for Phase III	Strong (to Good)

Part C. SOL Aquifer™ Fit Across the Six Categories

6. SOL Aquifer fit by category

Table 11 summarizes the fit in each category, with one worked example per category. Tables 12–17 then show, for every topic, the conditions SOL Aquifer would sweep, the physics solver that would generate the training data and what the surrogate would predict.

Table 11. SOL Aquifer fit by category (topic numbers)

Category	Worked example	Strong	Good	Partial
Weapon Systems & Lethality	01 Hypersonics: CFD and thermal-response runs; predicts heat load and thermal margin	01, 02, 03, 05, 06, 07, 08, 10	04, 09	None
Ground Vehicle Technologies	04 Vehicle performance: terramechanics and multibody dynamics; predicts speed made good and energy per km	04, 05, 06, 10	02, 03, 07, 09	01, 08
Networking & Communications	05 GNN network twins: packet-level simulation; a GNN predicts latency and deadline compliance	01, 02, 04, 05, 10	03, 07, 08	06, 09
Sensing Technologies	01 PNT sensor fusion: Monte Carlo INS and Kalman filter runs; predicts position error bound	01, 02, 06, 10	03, 04, 05, 07, 09	08
Air Platforms & Aviation	01 Future vertical lift: comprehensive rotorcraft analysis and CFD; predicts loads, fatigue margin, fuel burn	01, 03, 07, 08, 10	02, 04, 05, 06, 09	None
Soldier Performance	09 Heat and fatigue: thermoregulation model, the same heat balance as HVAC; predicts safe work time	01, 05, 06, 08, 09, 10	02, 04, 07	03

Across the 60 topics, SOL Aquifer fits best where a slow physics solver already exists. It works with a twist where the solver is a renderer that generates synthetic training data, where models are already fast and the value is calibration, or where hardware profiling replaces physics. It fits least where the key data comes from people, novel real-world events or formal proof. Aviation had no Partial topics because nearly every area has a physics solver. Soldier Performance Topic 09 was the closest analog of all, because a soldier's heat balance uses the same thermodynamics as that of a building, the case SOL Aquifer was first proven on.

Table 12. SOL Aquifer mapping: Weapon Systems & Lethality

#	Topic	Conditions swept	Physics solver	Surrogate predicts	Fit
01	Hypersonics	Mach, altitude, angle of attack, TPS material	CFD + thermal-response codes	Heat flux, surface temperature, thermal margin	Strong
02	Extreme-environment materials	Composition, processing and test temperature	Atomistic, thermodynamic and finite-element models	Strength, oxidation resistance	Strong
03	Material manufacturability	Laser power, scan speed, layer thickness	Additive-manufacturing thermal simulation	Porosity, residual stress, defect risk	Strong
04	Laser diagnostics	Temperature, pressure, species concentration	Spectroscopic line-by-line models	Gas state from a measured spectrum	Good
05	Directed energy	Range, turbulence, humidity, wavelength	Wave-optics beam propagation	Irradiance and beam quality	Strong
06	Fire control	Target motion, sensor noise, geometry	Tracking and engagement simulation	Track error and confidence	Strong
07	Air & missile defense	Threat profiles, sensor layout, comms loss	Sensor-network simulation	Track-fusion quality	Strong
08	Long-range precision fires	Wind profile, air density, launch conditions	6-DOF trajectory simulation	Flight path and dispersion	Strong
09	Energetics	Formulation, temperature, age	Thermochemical and aging models	Performance and stability indicators	Good
10	Cross-cutting T&E	Test parameters for any topic above	Any of the solvers above	Where to test next	Strong

Table 13. SOL Aquifer mapping: Ground Vehicle Technologies

#	Topic	Conditions swept	Physics solver	Surrogate predicts	Fit
01	One-to-many C2	Vehicles per operator, task tempo, comms delay	Discrete-event workload simulation	Operator workload and missed tasks	Partial
02	Off-road AI perception	Terrain, vegetation, lighting, weather	Physics-based camera and lidar simulation	Labeled synthetic scenes and expected accuracy	Good
03	Air-ground teaming	UAS altitude, route, terrain, occlusion	Multi-robot mapping simulation	Map coverage and route quality	Good (to Strong)
04	Vehicle performance	Soil, moisture, slope, speed, vehicle setup	Terramechanics + multibody dynamics	Traction, sinkage, speed made good	Strong
05	Power management	Mission profile, temperature, load, battery age	Electro-thermal battery and powertrain models	Range, battery temperature and health	Strong
06	Movement & maneuver	Terrain, route constraints, mission timing	Mobility and route simulation	Route time and risk	Strong
07	Long-duration ops	Duty cycle, temperature, component wear	Degradation and reliability models	Remaining useful life	Good
08	Human-AI collaboration	Explanation style, task difficulty	No physics solver; human studies	Operator trust and performance	Partial
09	Human-in-the-loop ML	Scenario difficulty, terrain, sensor quality	Autonomy simulation with scripted takeovers	Where the vehicle will need help	Good
10	Autonomy T&E	Scenario parameters	Full autonomy simulation stack	Failure regions and where to test next	Strong

Table 14. SOL Aquifer mapping: Networking & Communications

#	Topic	Conditions swept	Physics solver	Surrogate predicts	Fit
01	Resilient tactical networks	Mobility, jamming, terrain, node loss	Network emulation (ns-3, EMANE)	Route stability and delivery rate	Strong
02	AI/ML optimization	Spectrum load, interference, demand	Network and RF simulation	Throughput for a given allocation	Strong
03	Self-sensing networks	Fault type, jamming, terrain, hardware failure	Simulation with injected faults	Labeled degradation signatures	Good
04	Self-healing networks	Fault scenarios and candidate repairs	Emulation of each repair action	Outcome of a fix before it is applied	Strong
05	GNN network twins	Topology, traffic, link conditions	Packet-level network simulation	Latency, throughput, packet loss	Strong
06	Network security	Attack type, traffic mix	Attack traffic in a cyber range	Labeled attack traffic	Partial
07	RF spectrum awareness	Signal type, SNR, channel, multipath	RF channel simulation	Labeled synthetic signals	Good
08	Federated learning	Node count, bandwidth, dropouts, data skew	Federated learning simulation	Accuracy versus comms cost	Good
09	Secure data sharing	Policy set, device class	Crypto benchmarking on hardware	Latency and power per operation	Partial
10	Dual-use networks	Cell layout, user load, outage events	5G network simulation	Coverage and capacity	Strong

Table 15. SOL Aquifer mapping: Sensing Technologies

#	Topic	Conditions swept	Physics solver	Surrogate predicts	Fit
01	PNT sensor fusion	GPS outage length, IMU grade, dynamics, temperature	Monte Carlo INS + Kalman filter simulation	Position drift and integrity bound	Strong
02	Signals of opportunity	Emitter geometry, frequency, terrain, multipath	RF propagation and ray-tracing models	Ranging error and position accuracy	Strong
03	Vision-aided navigation	Season, lighting, altitude, terrain type	Physics-based image rendering	Map-matching accuracy and failure cases	Good (to Strong)
04	Precision timing	Temperature, vibration, holdover duration	Clock and oscillator noise models	Timing drift over holdover	Good
05	GPS interference	Jammer power, geometry, spoofing profile	GNSS signal and receiver simulation	Detection and false-alarm rates	Good
06	Collaborative PNT	Number of units, comms loss, geometry	Multi-agent navigation simulation	Team position error	Strong
07	Sparse-data ATR	Target, background, sensor, weather	Physics-based EO/IR scene simulation	Labeled synthetic imagery and expected accuracy	Strong
08	Multi-INT fusion	Sensor mix, noise, conflicting reports	Sensor and scenario simulation	Fused track accuracy	Partial (to Good)
09	Edge ISR processing	Model size, resolution, bandwidth, hardware	Hardware profiling runs	Latency, power and accuracy trade-off	Good
10	Novel sensors	Scene and sensor parameters	Sensor physics models (event camera, SAR)	Image quality and detection performance	Strong

Table 16. SOL Aquifer mapping: Air Platforms & Aviation

#	Topic	Conditions swept	Physics solver	Surrogate predicts	Fit
01	Future vertical lift Airspeed, altitude, gross	weight, maneuver	Comprehensive rotorcraft analysis + CFD	Hub loads, vibration, power required	Strong
02	Optionally manned	Scenario, weather, system faults	Flight simulation with autonomy in the loop	When autonomy should hand back control	Good (to Strong)
03	Degraded visual landing	Dust density, downwash, lighting, sensor type	Brownout particle and sensor simulation	Visibility and landing-aid accuracy	Strong
04	Rotorcraft health	Load spectrum, flight hours, seeded defects	Drivetrain dynamics and fatigue models	Fault vibration signatures, remaining life	Good
05	Swarm coordination	Swarm size, comms loss, wind, mission	Multi-agent flight simulation	Mission success and coverage	Good
06	Airborne sensors	Altitude, sensor type, weather, target	Sensor physics and scene rendering	Detection range and false alarms	Good
07	Materials	Layup, impact energy, temperature, erosion	Finite-element damage models	Damage extent, residual strength	Strong
08	Components	AM process settings, load spectrum	Process and fatigue simulation	Fatigue life	Strong
09	Runtime assurance	Flight envelope, gusts, controller faults	Closed-loop flight dynamics simulation	Where the monitor must switch to backup	Good
10	Hybrid-electric UAS	Mission profile, temperature, battery state	Electro-thermal powertrain models	Endurance and energy use	Strong

Table 17. SOL Aquifer mapping: Soldier Performance

#	Topic	Conditions swept	Physics solver	Surrogate predicts	Fit
01	Performance modeling	Workload, sleep, task demands	Biomathematical fatigue models	Alertness and performance over a mission	Good
02	LVC simulation	Terrain, force mix, scenario rules	Constructive simulation runs	Scenario difficulty and outcomes	Strong (to Good)
03	Adaptive training	Trainee skill, task difficulty	Simulated learners (knowledge tracing)	Expected learning gain	Partial
04	C2 decision aids	Scenario, information quality, time pressure	Wargame and constructive simulation	Outcome spread for each course of action	Good
05	Protective materials	Material, thickness, impact conditions	Finite-element impact models	Protection level versus weight	Strong
06	Soldier power	Mission profile, temperature, device load	Battery electro-thermal models	Runtime and battery health	Strong
07	Body-worn devices	Model size, sensor rate, hardware	Hardware profiling runs	Latency and battery life	Good
08	Physical performance	Load, gait, terrain, pace	Musculoskeletal simulation (e.g., OpenSim)	Joint loads and injury risk	Strong
09	Heat & fatigue	Temperature, humidity, sun, work rate, kit	Thermoregulation models	Core temperature and safe work time	Strong
10	First-responder use	Gear, environment, work rate	Same thermoregulation models	Safe work time for firefighters and EMS	Strong (to Good)

7. University mechatronics and additive manufacturing lab use

In every category, the university's mechatronics lab and additive manufacturing capability produce the real test data that SOL Aquifer's calibration step needs: the measured outcomes that drive the data flywheel. That also gives the university a concrete, hands-on share of the research. Figure 3 rates the value of the labs in each category, and Table 18 gives the detail.

Figure 3. Lab use across the six categories

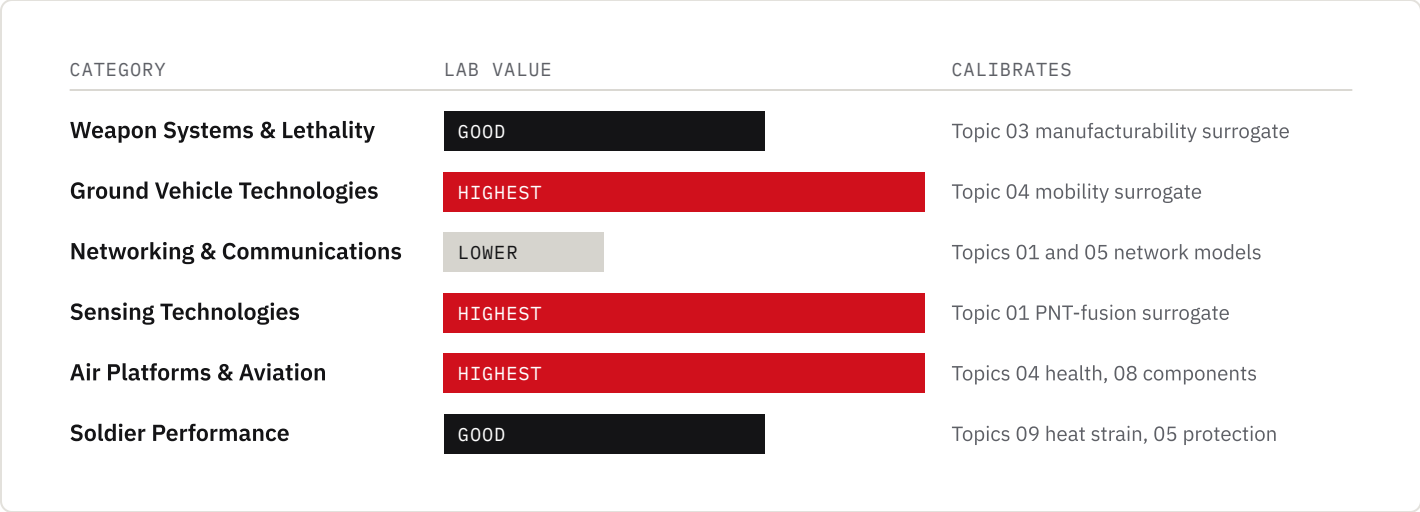

Lab value as rated in the report. The right column names the surrogate each category's lab data calibrates. Redrawn from the report figure.

Table 18. Lab use by category

Category	Value	Mechatronics	Additive manufacturing	Calibrates
Weapon Systems & Lethality	Good	Pointing and tracking testbeds (gimbals, fast steering mirrors); edge hardware for laser diagnostics	Monitored, CT-scanned test coupons give labeled defect data; subscale fixtures; ceramics if available	Topic 03 manufacturability surrogate
Ground Vehicle Technologies	Highest	Single-wheel soil-bin rig; 1/10-scale robot testbeds for autonomy and takeovers; battery HIL bench	Quick-turn wheel and track designs for the soil bin; sensor mounts, chassis parts	Topic 04 mobility surrogate
Networking & Communications	Lower	Mobile robots carrying software-defined radios as a mesh testbed; shielded-enclosure jamming tests	Printed antenna prototypes, RF enclosures, radio mounts	Topic 01 and 05 network models
Sensing Technologies	Highest	Motion platforms, turntables and robot arms with motion capture for ground-truth IMU and camera data	Precisely aligned multi-sensor rigs, housings, calibration targets	Topic 01 PNT-fusion surrogate
Air Platforms & Aviation	Highest	Rotor test stand with seeded gear and bearing defects; indoor drone flight cage; hybrid-electric dyno	Subscale rotor blades and drone airframes; fatigue coupons for AM components	Topic 04 health and Topic 08 component surrogates
Soldier Performance	Good	Instrumented thermal manikin or heated forms (fewer human trials); wearable and exoskeleton prototypes	Printed lattice padding for helmet and armor liners; custom-fit wearable housings	Topic 09 heat-strain and Topic 05 protection surrogates

Confirm with the university:

- **Equipment inventory.** Polymer versus metal or ceramic additive manufacturing capability. Polymer AM covers fixtures, housings and robots; materials and component topics need metal or ceramic AM.
- **Controlled space.** Directed energy, hypersonics and aviation components may be ITAR controlled, needing restricted lab access and limits on foreign-national students.
- **Budget line.** Lab time and materials budgeted as part of the university's share of the work, not as overhead.

References

1. SOL Engineering Services (2026). *Modeling HVAC System Performance: Code as Architecture Approach to AI*. Poster Summary, Vol. 2.
2. U.S. Army DEVCOM Army Research Laboratory. ARL-BAA-0038, *Artificial Intelligence and Machine Learning with Extremely Sparse Data*.
3. U.S. Army DEVCOM Army Research Laboratory. ARL-BAA-0107, *Academic Applied and Advanced Innovation*.
4. U.S. Army DEVCOM Army Research Laboratory. ARL-BAA-0064, *Autonomous Sensing and Information Fusion for Advanced Indications and Warnings*.