Scout-and-Maneuver: A GNSS-Free Air-to-Ground 3D Spatial Frame for Ground-Vehicle Navigation in Contested Terrain

Guillaume Cortes, Yannick Menet, Lionel Bouchard — Tyndall FX, France
Contact: guillaume.cortes@tyndallfx.com

Identification of Use Cases

Operational context. A ground vehicle, crewed or robotic, sees its surroundings from ground level, with ridgelines, vegetation and buildings blocking its line of sight. Two things make that harder today: terrain that is unmapped or has recently changed — craters, rubble, new obstacles — and the routine denial of GNSS by jamming. The crew or the autonomy stack is then left without a current map and without satellite positioning at the moment it has to choose a route, and a bad route costs exposure or immobilisation.

The concept. Our approach is to put the sensor in the air. A UAV scout flies ahead of the vehicle and reconstructs the terrain the vehicle cannot yet see, in georeferenced 3D and in real time. It also establishes a common spatial frame that both platforms share without GNSS, and that is the contribution we want to put forward here. The ground vehicle localizes itself and plans its route inside that freshly built model. All of this runs on the unit's own compute: scout, model and vehicle form a self-contained loop at the tactical edge. We have demonstrated this UAV–UGV chain with French defence stakeholders, using the scout's 3D reconstruction as the input to a ground robot's autonomous navigation.

Intended users. We are aiming at armoured and robotic ground units working reconnaissance, convoy movement, route clearance and force protection. The same chain transposes to civil security, in particular disaster response and search-and-rescue under GNSS denial.

Representative use cases.

1. **Route reconnaissance and path-planning ahead of a ground vehicle.** The scout maps the axis of advance a few minutes before the vehicle gets there, leaving the crew or the autonomy stack time to re-plan around a route that turns out to be blocked or exposed.
2. **Manoeuvre under GNSS denial.** When GNSS is jammed or otherwise denied, the vehicle continues to navigate on the shared frame built from the air, and keeps moving where a GNSS-dependent vehicle would have to stop.
3. **Operations in unmapped or just-changed terrain.** Where no current map exists — expeditionary, urban or post-strike settings — the scout produces a fresh 3D picture on demand and updates it as the situation changes.

Use-case benefit. Manual route reconnaissance is slow and it exposes the people who carry it out; this takes over part of that work. It also keeps the force moving when GNSS-guided movement would have to stop.

Proposed Technical Approach

Proposed AI approach. We propose an air-to-ground spatial-framing loop (Figure 1) that runs on the unit's own compute and needs no server, no cloud connection, no GNSS during operation and no pre-existing map. It combines two Tyndall FX modules — a 3D reconstruction module (Terramind) and a localization module (Geomind) — with two further functions that tie them into a single cross-platform frame:

1. **Real-time 3D reconstruction with no prior map.** The 3D reconstruction module turns the scout's imagery into a georeferenced metric 3D model of the terrain in real time, with no map or scan to start from. It recovers oblique and façade geometry, and works from thermal or night-vision imagery as well as daylight, so the loop still functions at night and in poor visibility.
2. **One GNSS-free frame across two moving platforms.** The localization module estimates 6-DoF pose from visual odometry and an IMU, with no RF or satellite signal, so the scout and the vehicle can share one georeferenced frame while GNSS is jammed. A single initial position fix is needed at start-up — taken from GNSS, or entered by the operator — and nothing after that. The shared model then acts as the reference: each platform re-anchors against it continuously, which is what keeps the error from growing with the distance travelled. On the aerial platform we measure drift below 1 % of the distance flown, with no satellite signal once the flight has begun. The module reports the position and attitude estimate with its uncertainty; navigation and guidance stay with each platform's own autopilot or autonomy stack.
3. **Air-to-ground registration.** The model built from the air is registered into the ground vehicle's own viewpoint and frame. This is a cross-view alignment between an overhead reconstruction and a ground-level sensor stream. Solving it is what makes the model usable from where the vehicle stands.
4. **Traversability and threat layer.** Semantic classification of the shared model produces a mobility layer — slope, obstacles, cover and concealment, choke points — georeferenced in the same frame. Where the aerial view is occluded, the layer marks the area as unobserved instead of filling it in. AI detection adds the second half of the picture: people and vehicles found in the scout's imagery are anchored in the 3D frame, so a detection along the axis of advance carries a position and not only a box in an image. Both layers are meant to be read by the vehicle's own navigation tools, which can then plan a trajectory that accounts for them.

Where the computation runs. The reconstruction can run at three complementary points, chosen per mission depending on whether an operator is in the loop and how much bandwidth the datalink leaves. **On the operator's tablet or laptop:** the unit sends a ground robot forward, flies a scout ahead of it, builds the model from the downlinked feed and hands it to

the robot already deployed. **On the scout, in the air:** the model is built in flight and sent down to the robot as a compact georeferenced model instead of a video stream, at a fraction of the bandwidth. **On the robot:** the video comes down and the robot builds and consumes the model on its own compute, with no operator in between. The robot also carries and recovers the scout, which gives the pair a way to work when electronic warfare closes the air-to-ground link altogether: the scout lands back on the robot, and the robot reads the data off it directly, already processed or still to be processed. This path was part of the design from the start, so losing the link costs the crew time rather than the terrain picture itself.

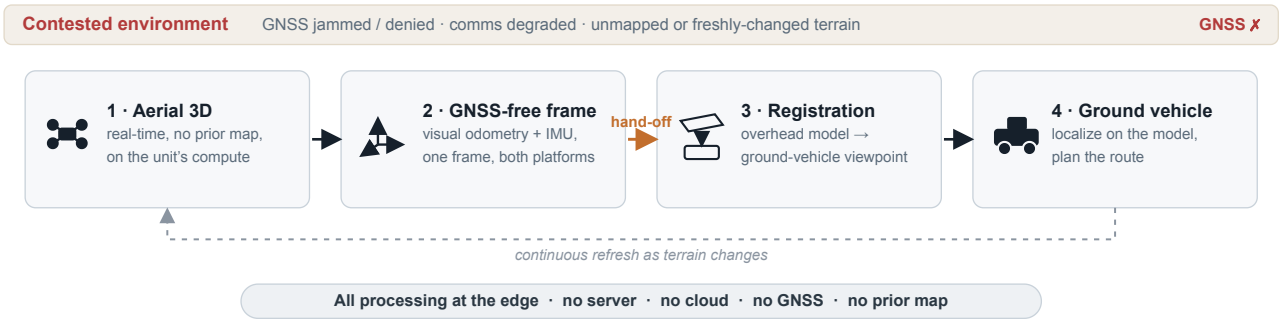

Figure 1. The scout-to-maneuver loop, left to right: (1) the UAV scout's imagery is reconstructed into real-time 3D with no prior map; (2) both moving platforms share one georeferenced frame without GNSS; (3) the model is registered into the ground vehicle's viewpoint; (4) the vehicle localizes itself and plans its route on the model. The loop refreshes as the terrain changes, and every stage runs at the edge — on the tablet, the robot or the air vehicle. *(Illustrative synthetic schematic.)*

What is new. Cloud-based photogrammetry and GNSS-dependent UAV mapping either build their model somewhere else, or need satellite positioning to build it at all; here the model is built on the spot, in real time, without either. Offline SLAM pipelines produce a comparable model, but only once the mission is over. The comparison that matters most, though, is with reference-map matching, the approach usually taken to absolute positioning without GNSS. What separates the two is the age of the terrain data. A 2D satellite basemap is typically months or years old and flat, while the scout's model is a few minutes old and three-dimensional. Where a strike or a collapse has changed the ground, the stored reference no longer describes what is there, and map-matching loses accuracy in exactly the situations the mission depends on. The scout builds its own reference instead of consuming one, and the pair navigates on that. The approach extends CAID 2025's WarNav [1] (navigable-zone segmentation within a single frame) by carrying a metric georeferenced frame across two independently moving platforms, and complements ORCA [2] (air-land manned-unmanned C2 teaming) by contributing the spatial substrate rather than the command layer. Our main technical contribution is the combination of the GNSS-free cross-platform frame with air-to-ground registration; neither subsystem on its own is the difficult part.

Technical challenges. Most of what remains difficult sits in the machine-to-machine step: registering an overhead model into a ground viewpoint, across differences in scale, perspective and appearance; holding a single GNSS-free frame across two platforms that move independently; producing a semantic and detection layer that a machine can act on with no human to interpret it; and fitting the whole chain inside the bandwidth a contested datalink actually leaves. Building the reference rather than reading one also costs us coverage: a scout maps hundreds of metres, not a theatre of operations, so we gain freshness and give up extent. Real-time 3D from a moving airborne sensor, with no prior map and inside edge compute and latency budgets, is already implemented and running — on commodity tablet-class hardware, not server-class GPUs.

Maturity level. The human-in-the-loop chain has been demonstrated end-to-end: a scout's reconstruction, built on the operator's tablet or laptop, handed to a ground robot that then navigates on it. Real-time 3D reconstruction, GNSS-free localization and AI detection anchored in the model are mature building blocks. The machine-to-machine elements remain at prototype maturity and are under active development — the semantic traversability layer the robot reads to adapt its own navigation, the air-to-ground registration, and the land-and-download recovery path. Those are the focus of the coming experimentation campaigns. Reconstruction hosted on the scout in flight, and fusion across several scouts, are on the roadmap. Over a 200 × 200 m area the scout's flight takes under 90 seconds, the reconstruction is produced in real time, and obstacle dimensions are measured to within 2 %, validated in operational conditions with customers and partners, both against reference maps and by measuring objects of known dimensions. That level of accuracy matters because obstacles have to be sized correctly to be avoided reliably, and missing them is a common cause of failure for autonomous ground mobility and for manoeuvre planning.

Deployment perspectives. Deployment adds little to what a unit already carries: the loop runs on assets it already fields, with nothing external to stand up or connect to. Integration is towards existing vehicle autonomy stacks and C2/GIS systems, through the shared georeferenced model. Our priority next step is the machine-to-machine pairing with the model built on the robot, planned with French defence stakeholders over the coming months. Field experimentation with operational units follows, to test the frame and the layers against real terrain and real mission conditions, and to establish where the system can be relied on.

References

1. M.-E. Coupvent des Graviers, H. Ammar, C. Guettier, Y. Dumortier, R. Audigier, "WarNav: An Autonomous Driving Benchmark for Segmentation of Navigable Zones in War Scenes," *CAID 2025*, Safran Electronics & Defense / Université Paris-Saclay, CEA List.
2. A. Amrani, J.-B. Blanc–Rouchossé, L. Boudet *et al.*, "ORCA, a Trustworthy Symbolic Artificial Intelligence Assistance in Air-Land Collaborative Combat Combining Manned and Unmanned Forces," *CAID 2025*, Université Paris-Saclay CEA List / TNS / Thales.