# Research Basis for a Bounded Microfarm Rover

## Executive summary

The premise is defensible: a microfarm rover is technically and economically viable when it is bounded to lightweight, low-speed, modular, supervised semi-autonomous work on plots of roughly 0.2–2 hectares, rather than framed as a miniature general-purpose tractor or a fully unsupervised farm AI. That framing is consistent with the uploaded design documents, which converge on an electric, modular rover with a nominal 48V 30Ah LiFePO4 battery, a 75–150 kg platform mass centered around 120 kg, about 80 kg payload, narrow-row operation, and an explainable autonomy layer based on machine perception plus rules. Externally, that design logic matches the mechanization gap identified by FAO[1], the adoption constraints identified by GAO[2], and the current safety-oriented trajectory of agricultural autonomy standards. [3]

The most rigorous technical claim is not "full autonomy." It is that a supervised semi-autonomous rover using RTK-GNSS, IMU, wheel odometry, row-relative vision, conservative local planning, independent obstacle protection, and a deterministic MCU safety layer can reliably perform a first wave of high-value tasks: precision seeding, granular micro-dosing, scouting, irrigation support, and mechanical weeding. ROS[4] 2 and Nav2 already support the necessary outdoor-navigation building blocks, including GPS-based localization workflows, lifecycle-managed nodes, and a collision-monitor layer that sits outside the nominal planner. Meanwhile, field literature shows that row vision, GNSS fusion, and targeted machine vision for crop and weed identification are mature enough for these specific tasks, while dexterous harvesting and unconstrained spraying remain much higher risk. [5]

The economic case is strongest when utilization is shared. Single-owner economics can work in high-value urban or peri-urban vegetable systems, but the strongest generalizable model for lower-margin microfarms is cooperative ownership or Rover-as-a-Service. That conclusion follows directly from the FAO hire-service literature, the broader smallholder mechanization work associated with CGIAR[6], and the GAO finding that precision technologies can reduce inputs and improve efficiency but are constrained by cost, interoperability, and data issues. A rover that is modular, repairable, and used at high annual utilization can clear that barrier; one that sits idle on a single low-margin plot often will not. [7]

## Scope, assumptions, and evidence base

This report integrates the uploaded documents as the design brief and tests them against primary and near-primary sources: standards, official documentation, manufacturer datasheets, public-agency reports, field-robot case material, and recent peer-reviewed literature. The uploaded documents are internally consistent on the rover's general character: small, electric, modular, locally repairable, slow, explainable, and farmersupervised. The report therefore treats the documents as the requirement baseline and evaluates whether that baseline is technically and economically supportable.

The explicit assumptions below are either user-specified or directly implied by the uploaded design materials. They should appear verbatim in any paper draft so that reviewers can distinguish claimed evidence from scenario modeling.

| Assumption | Working baseline |
|---|---|
| Target farm size per rover | 0.2–2 ha |
| Platform mass | 75–150 kg, nominal 120 kg |
| Payload | 80 kg |
| Width | 0.8–1.2 m |
| Battery | 48V 30Ah unless otherwise noted |
| Battery chemistry | LiFePO4 baseline; NMC analyzed as alternative |
| Autonomy | Supervised semi-autonomy |
| Work speed | ~0.5–2 km/h in-task |
| MVP modules | Seeder, fertilizer micro-doser, scouting, mechanical weeder |
| Deferred modules | Full autonomous spraying, dexterous harvesting |

Those assumptions are credible for smallholder and market-garden mechanization because FAO's sustainable mechanization guidance emphasizes appropriateness of scale, business-model fit, hire services, and equipment matched to local operating realities rather than miniaturized large-farm machinery. GAO likewise finds that precision agriculture's benefits come from more exact delivery of inputs and more timely, data-driven decisions, while its barriers are cost, interoperability, and usability. A microfarm rover should therefore be argued as a scale-matched precision implement carrier, not as a compact tractor substitute. [8]

## Technical design and autonomy

A rigorous paper should recommend a four-wheel, high-clearance, low-center-of-gravity rover with swappable modules. Wheels should be preferred over tracks for the baseline system because they are cheaper, easier to source and repair, lower-drag on firm beds and row alleys, and align with the uploaded documents' emphasis on commodity components and local repairability. Tracks are better treated as a niche option for greenhouse floors or persistently soft soils, not the primary architecture. Field-robot soil literature also supports the basic mass logic: lighter machines reduce deep-compaction risk relative to heavy tractors, but repeated passes still matter, especially in topsoil, so traffic planning and dedicated lanes remain important. [9]

A credible drive train for the baseline rover is either four geared wheel motors with encoder feedback or two driven wheels plus two passive/support wheels if cost dominates and terrain is modest. Steering can be skid-steer or differential drive at MVP, but row crops and narrow beds generally benefit from conservative turning behaviors and headland routines rather than aggressive pivoting. The uploaded documents are right to emphasize slow work rates. Below roughly 2 km/h, perception error and tool-actuation latency become much easier to manage, which is exactly where this rover should live.

Representative current component choices are summarized below. The point is not that there is one inevitable bill of materials. The point is that the market now contains a credible low-cost stack and a credible ruggedized stack for each critical function. Component prices below are indicative current listed prices and will move with memory and electronics markets. [10]

| Subsystem | Lean option | Higher-performance option | Research judgment |
|---|---|---|---|
| Compute | Raspberry Pi[11] Pi 5 4GB | NVIDIA[12] Jetson Orin Nano Super | Pi-class compute is enough for deterministic autonomy and light CV; Jetson-class compute is justified if the rover runs multiple neural detectors, VLM/VLA captioning, or edge analytics |
| RTK GNSS | u-blox[13] ZED-F9P breakout class | Emlid[14] Reach M2 / RS2+ class | ZED-F9P is a strong cost/performance baseline; integrated GNSS receivers make sense when deployment simplicity and field ruggedness matter more than BOM minimum |
| Depth / row vision | Luxonis[15] OAK-D Lite | Intel[16] RealSense D455 or industrial stereo | OAK-D Lite is the best lean edge-vision option; D455-class hardware adds range and maturity for obstacle perception |
| 2D LiDAR | SLAMTEC[17] RPLIDAR A2 class | Hokuyo[18] UST-10LX class | Use LiDAR when dust, low light, or safety margin demands it; 2D LiDAR is more defensible as a safety aid than as the only localization sensor in crops |
| Battery | Bioenno Power[19] 48V 30Ah LiFePO4 class | Larger LiFePO4 pack or swappable dual-pack architecture | LiFePO4 is the right default chemistry for the rover baseline |
| Solar charging | Victron Energy[20] SmartSolar MPPT 100/20 class | Larger MPPT and docked solar array | MPPT is recommended over PWM for both dock charging and on-rover solar assistance |

| Subsystem | Lean option | Higher-performance option | Research judgment |
|---|---|---|---|
| Traction motors | Pololu[21] 37D encoder gearmotor class for prototype | Teknic[22] integrated servo or Roboteq[23] industrial controller tier | Commodity gearmotors are fine for MVP if torque margins are conservative; industrial servo/control hardware matters for repeated daily commercial duty |

Source note: current listed reference points include Jetson Orin Nano Super at $249 and 67 TOPS; Raspberry Pi 5 4GB around $85 at DigiKey; SparkFun ZED-F9P breakout at $259.95; OAK-D Lite at $169; RealSense D455 at $419; Reach M2 at $749; Reach RS2+ at $2399; Bioenno 48V 30Ah LiFePO4 at $549.99; Victron MPPT 100/20 supporting up to 48V battery systems and up to 1,160 W nominal PV at 48V; Pololu 24V 37D gearmotors from about $60.95 with encoder variants; and Roboteq safety-capable BLDC controller tiers beginning in the hundreds of dollars. [24]

Localization should be argued as GNSS-first, fusion-based, and crop-relative locally. In open fields, the most defensible stack is RTK-GNSS plus IMU plus wheel odometry, fused through robot_localization or an equivalent estimator, with row-relative vision used to correct local alignment and reduce crop damage. Nav2’s GPS localization workflow explicitly supports this architecture, and the u-blox ZED-F9P class of receivers is designed for centimeter-level positioning with RTK and moving-base support. Recent agricultural navigation papers also support multi-sensor fusion and row-oriented visual perception as the practical path in dynamic crop environments. [25]

The paper should be careful with SLAM. SLAM is situational, not foundational, for this rover. In GNSS-friendly open plots, full LiDAR-SLAM or visual SLAM should not be sold as the primary positioning backbone. It becomes important in orchards, tunnels, greenhouses, or heavily degraded GNSS conditions. Recent papers show that agricultural SLAM is progressing, but it still faces lighting variability, repetitive geometry, vegetation motion, dust, and seasonal appearance changes. For the microfarm rover, reviewers are more likely to trust “RTK fusion with optional SLAM fallback” than “SLAM-centric field autonomy.” [26]

Obstacle avoidance belongs in a layered safety architecture. Planner-level costmaps are necessary but not sufficient. A separate collision-monitor path is preferable because it can issue immediate slowdown or stop commands on short notice, independent of the nominal path planner. That is exactly how Nav2’s collision monitor is described, and it aligns with the newer ISO 18497 partitioning between machine design, obstacle protection, operating zones, and verification/validation. [27]

```
flowchart LR
    subgraph Perception
        A[RTK-GNSS]
        B[IMU + encoders]
```

```
        C[Row camera / depth camera]
        D[Optional 2D LiDAR]
        E[Soil, hopper, current, moisture sensors]
    end

    subgraph Edge stack
        F[MCU safety layer]
        G[ROS 2 nodes]
        H[State estimation + Nav2]
        I[VLM/VLA captioning]
        J[Safety and agronomy rules]
    end

    subgraph Actuation
        K[Traction motors]
        L[Tool lift / linear actuators]
        M[Seeder / doser / weeder / pump]
    end

    A --> H
    B --> H
    C --> H
    D --> H
    C --> I
    E --> J
    H --> J
    I --> J
    J --> F
    F --> K
    F --> L
    F --> M
```

The uploaded documents’ idea of a captioning layer is strongest when treated as an explainability and exception-management tool, not as the only controller. RT-2 and OpenVLA show that vision-language-action models can transfer semantic knowledge into robotics and can be fine-tuned efficiently, but that is not the same thing as giving a foundation model direct authority over field tools. The rigorous argument is narrower: use a VLM/VLA layer to generate human-readable field understanding, detect anomalies, summarize why the rover paused, and support operator approval workflows; keep motion and tool authority in the deterministic stack. [28]

```
flowchart TB
    GNSS[RTK GNSS mast]
    FRONT[Front RGB/depth camera]
    LIDAR[Optional forward LiDAR]
    IMU[IMU near center of gravity]
    BODY[[Rover chassis]]
    ENC[Wheel encoders]
    TOOL[Rear or underbody tool camera]
```

```
    HOP[Hopper and flow sensors]

    GNSS --> BODY
    FRONT --> BODY
    LIDAR --> BODY
    IMU --> BODY
    ENC --> BODY
    TOOL --> BODY
    HOP --> BODY
```

## Power, payload, and software

The baseline battery chemistry should be LiFePO4, not NMC, for the first commercial rover. The overall chemistry trade-off is straightforward. The IEC[29] and battery market literature support LFP/LiFePO4 for applications that prioritize cycle life, thermal stability, and safety margin; the IEA[30] reports that LFP packs are roughly 30% cheaper per kWh than NMC packs, while NMC still maintains a mass-based energy-density advantage of roughly one-fifth and a volumetric advantage of about one-third. For a low-speed rover with modest daily range and high safety sensitivity, that is a favorable trade-off for LiFePO4. [31]

With the stated baseline battery, nominal stored energy is 1.44 kWh. Using a conservative routine-use basis of about 80% usable energy gives roughly 1.15 kWh available per cycle. That is enough for a working-day rover if the toolset stays in the precision-assist class. It is not enough to turn the machine into a high-draft tillage platform. This is where the uploaded documents are technically sound: they repeatedly preference shallow strip opening, seed placement, scouting, and micro-dosing over full-width heavy tillage.

The power architecture should explicitly compare overnight AC charging, midday battery swap, on-rover solar canopy, and dock-based solar with MPPT. A paper that presents solar as “free propulsion” will look sloppy; a paper that presents solar as an energy extender and resilience layer will look rigorous. FarmDroid is the strongest published market proof point here: its FD20 uses four solar panels, high-precision RTK seeding/weeding, about 1.6 kWh of battery output, and 18–24 hours of operation within its own large-frame design envelope. That proves the concept of solar-powered agricultural robotics. It does not prove that every smaller rover should be solar-only. FarmDroid[32] shows that solar-first field robotics is possible, but it does so on a much larger machine with relatively large panel area and tightly optimized task scope. [33]

A practical energy budget for the 48V 30Ah rover is shown below.

| Mission profile | Average electrical load | Usable energy basis | Modeled runtime |
|---|---:|---:|---:|
| Scouting / mapping / logging | 120 W | 1.15 kWh | ~9.6 h |
| Precision seeding | 160 W | 1.15 kWh | ~7.2 h |
| Seeding + fertilizer micro-dosing | 190 W | 1.15 kWh | ~6.1 h |

| Mission profile | Average electrical load | Usable energy basis | Modeled runtime |
|---|---|---|---|
| Mechanical weeding / irrigation support | 220 W | 1.15 kWh | ~5.2 h |
| Shallow strip opening / aggressive cultivation | 350 W | 1.15 kWh | ~3.3 h |

These are modeled engineering estimates, not manufacturer duty-cycle guarantees. Their purpose is analytical: they show that the platform is strongest for low-speed precision operations, exactly as GAO's precision-agriculture framing would imply. [34]

For solar charging, MPPT is not optional if the project includes a solar dock or canopy. Victron's SmartSolar MPPT line is a useful conservative reference: the 100/20 variant supports 48V battery systems, nominal PV power up to 1,160 W at 48V, and peak efficiency around 98%, which is more than sufficient for a dock charger or modest canopy-assisted architecture. [35]

```
pie showData
    title Example average power split during precision seeding
    "Traction" : 55
    "Tool actuation" : 50
    "Compute + perception" : 30
    "GNSS + comms + IO" : 10
    "Conversion losses + margin" : 15
```

Payload and modularity should be sequenced by technical maturity and regulatory burden. The paper will be strongest if the modules are staged this way:

| Module | Technical maturity | Regulatory / safety burden | Recommendation |
|---|---|---|---|
| Precision seeder | High | Low | Launch module |
| Fertilizer micro-doser | High | Low | Launch module |
| Mechanical weeder | Medium-high | Medium | Phase-two module |
| Irrigation assist | Medium | Low-medium | Phase-two module |
| Spot sprayer | Medium | High | Later, human-approved only |
| Harvest assist logistics | Medium | Low-medium | Later add-on |
| Dexterous autonomous harvester | Low for MVP | Medium-high | Research only in first paper |

This ordering is well supported. Mechanical weeding is now well enough studied to justify inclusion in a serious roadmap, but fully automated picking remains constrained by perception complexity, occlusion, crop damage, and end-effector difficulty. Recent harvesting reviews make that point clearly, while weed-detection and real-time plant/weed recognition papers show that perception for precision weeding is much closer to deployment. [36]

The software architecture should cleanly separate hard-real-time and soft-real-time concerns. The lower layer should be an MCU-based safety and control subsystem handling motor loops, relays, E-stop evaluation, watchdogs, actuator limits, and fault states. The upper layer should run ROS 2, robot_localization, Nav2, mission logic, and higher-level perception. This separation is not cosmetic. It is exactly what makes a regulator, insurer, or safety reviewer more likely to trust the design. ROS 1 Noetic went end-of-life on 2025-05-31, so a new rover paper should not recommend ROS 1 as its strategic software base. [37]

A research-grade stack therefore looks like this:

| Layer | Recommended role |
|---|---|
| MCU safety layer | E-stop, watchdogs, interlocks, safe states, low-level IO |
| Motor / actuator control | Wheel velocity control, lift control, hopper and pump control |
| ROS 2 middleware | Drivers, transforms, timing, topics/services/actions |
| State estimation | GNSS + IMU + wheel encoder fusion, crop-row local correction |
| Nav2 mission layer | Waypoints, row missions, behavior trees, costmaps, collision monitor |
| ML perception | Row detection, crop/weed detection, plant-state and anomaly detection |
| VLM/VLA layer | Captioning, operator explanations, exception summaries |
| Data pipeline | Local-first logging, deferred sync, prescription ingestion, task audit trail |

For ML models, the paper should recommend task-specific detectors and segmentation models first, typically YOLO-family or equivalent compact detectors for crop/weed localization, with disease or pest classifiers optionally running on sampled scouting images. The best-supported design is “small, explainable, edge-deployable models first; large multimodal reasoning second.” Recent work shows strong field performance from advanced YOLO variants for weed/crop detection and credible edge-oriented pest/disease inference from ground robots. [38]

## Agronomy and deployment workflows

Agronomically, the rover is best matched to row crops and bed-based market-garden crops where plant spacing, travel corridors, and intervention timing can be prescribed. Good first-use crops include maize, beans, sorghum, onions, brassicas, leafy greens, peppers, and tomatoes, depending on bed geometry and canopy height. This also aligns with the uploaded documents, which emphasize maize, beans, sorghum, and structured narrow-row operations.

The strongest agronomic argument is timeliness plus placement. FAO's mechanization work emphasizes the productivity benefit of timely operations, and GAO's precision-agriculture review emphasizes that exact application of inputs such as water and fertilizer can improve efficiency and reduce waste. A rover that places seed more consistently, doses fertilizer near the root zone, and repeats scouting at high frequency is agronomically valuable even before it becomes highly autonomous. [39]

Soil interaction must be treated conservatively. The rover should be positioned for shallow strip opening, seed-slot creation, local cultivation, and topsoil disturbance control, not for full-width inversion tillage. That position is well aligned with both the uploaded documents and soil-compaction research. Lightweight robots reduce stress compared with large machines, but repeated passes still compact the surface if traffic is unmanaged, so the rover should be designed around permanent lanes where possible, careful wheel selection, and mission planning that minimizes excess traversal. [9]

Microclimate management should also be stated carefully. The rover does not "control climate." What it can do is support soil-moisture-informed irrigation, canopy-temperature scouting, humidity and temperature logging, and repeated spatial sensing. EPA and USDA-related irrigation guidance show that soil-moisture-based scheduling can reduce water waste and improve timing of irrigation decisions, which supports the inclusion of moisture probes and irrigation-assist modules in the rover roadmap. [40]

The operational workflow is strongest when it is explicit, auditable, and geofenced.

```
flowchart TD
    A[Field setup and geofence] --> B[Preflight checks]
    B --> C[Select module and agronomy prescription]
    C --> D[Localization and row alignment]
    D --> E[Run row mission]
    E --> F[Continuous obstacle and fault monitoring]
    F --> G{Exception?}
    G -- No --> H[Continue mission]
    G -- Yes --> I[Pause, alert, caption explanation]
    I --> J[Operator approve, adjust, or abort]
    J --> H
    H --> K[Dock or battery swap]
    K --> L[Sync logs and task records]
    L --> M[Maintenance and calibration]
```

Deployment differs materially by context.

| Scenario | Best-fit tasks | Best power mode | Best autonomy mode | Best ownership model |
|---|---|---|---|---|
| Urban microfarm | Seeding, scouting, irrigation support, logistics | Overnight AC + optional canopy | Tight geofence, close supervision | Direct ownership or neighborhood share |

| Scenario | Best-fit tasks | Best power mode | Best autonomy mode | Best ownership model |
|---|---|---|---|---|
| | assist | | | |
| Peri-urban vegetable farm | Seeding, weeding, scouting, dosing | AC + swap battery | Supervised row autonomy | Direct ownership or co-op |
| Rural microfarm / smallholder | Seeding, fertilizer micro-dosing, scouting | Solar dock + service routing | Geofenced supervised autonomy | Co-op or service-provider model |

FAO’s hire-service work strongly supports this deployment logic. Small farms rarely need full-time ownership of every machine, but they often benefit from access to timely service when planting windows matter. That same reasoning carries over to a rover platform. [41]

## Economics and business models

The economics should be presented in two layers: a hardware-only MVP build estimate and a modeled ownership/payback analysis.

A realistic prototype-grade MVP BOM for the baseline rover looks like this:

| Cost block | Modeled range |
|---|---:|
| Frame, chassis fabrication, wheels, bearings | \$900–\$1,500 |
| Drive motors and controllers | \$1,000–\$2,200 |
| 48V 30Ah LiFePO4 battery + charger | \$600–\$800 |
| Solar canopy or dock-side solar electronics | \$350–\$900 |
| RTK GNSS | \$260–\$750 |
| Cameras / perception sensors | \$170–\$500 |
| LiDAR optional add | \$150–\$1,200 |
| Compute + MCU + safety IO | \$150–\$500 |
| Wiring, connectors, enclosures, relays, E-stop | \$500–\$1,000 |
| Seeder + micro-doser launch module | \$700–\$1,500 |
| Assembly contingency | \$800–\$1,200 |
| Prototype-grade total | ~\$5,580–\$10,050 |

A fair interpretation is that a serious hardware MVP is likely to land around \$6.5k–\$8.5k, while a more rugged commercial-ready alpha unit will likely land nearer \$9k–\$12k+ before certification, NRE, and support overhead. That range is consistent with the current listed prices on the representative parts cited earlier plus modeled fabrication and assembly costs. [42]

Operating costs should be modest but not ignored. A defensible working model for annual opex per rover is $700–$1,500, covering maintenance consumables, replacement wear parts, battery reserve, connectivity/logging cost, and seasonal calibration effort. That is precisely why utilization matters so much: a lightly used rover does not amortize its fixed capex effectively.

The most rigorous economic claim is that direct ownership is context-dependent, but shared-utilization models are broadly more robust. FAO's training material for mechanization hire services is explicit that small-scale service provision can be a viable business enterprise when equipment choice, maintenance, local demand, and business management are aligned. CGIAR/CIMMYT work in southern Africa similarly emphasizes that the service-provider model succeeds only when machinery scale, training, after-sales service, and local business conditions fit the environment. [41]

Illustrative ROI scenarios under the stated assumptions are shown below. These are modeled scenarios, not observed market averages.

| Scenario | Capex basis | Annual net benefit assumption | Modeled payback |
|---|---:|---:|---:|
| Urban direct owner, high-value crops | $8,500 | $2,500–$4,000 | 2.1–3.4 years |
| Peri-urban vegetable farm | $9,500 | $2,200–$3,600 | 2.6–4.3 years |
| Rural single microfarm, staple-heavy mix | $8,500 | $800–$1,800 | 4.7–10.6 years |
| Five-farm cooperative | $10,500 | $4,500–$7,500 | 1.4–2.3 years |
| Rover-as-a-Service operator | $10,500 | $5,000–$8,000 contribution margin | 1.3–2.1 years |

The direction of this table is the important finding. High-value vegetables and shared-utilization models work best. A single low-margin staple microfarm may still benefit operationally, but may struggle to justify direct ownership financially. That is fully consistent with the FAO and GAO evidence base. [43]

For the paper's final figure set, the two most informative additional visuals would be a stacked bar chart of subsystem power by task and a utilization-versus-payback sensitivity curve. Those figures would make the central economic argument much clearer than a single ROI table.

## Safety, lifecycle, and environmental impact

A rigorous paper must anchor the rover in current safety and regulatory frameworks. The core machinery-safety scaffold is ISO[44] 18497, which in its 2024 split now separately addresses machine design principles and vocabulary, obstacle protection, autonomous operating zones, and verification/validation. That structure maps directly onto a geofenced microfarm rover. For safety-related control systems, ISO 25119 remains the relevant series for agricultural machinery SRP/CS design and development. A credible paper should therefore state that the rover's hazard controls are being designed against ISO 18497 parts 1–4 and ISO 25119 parts 1–4. [45]

In the United States, chemical modules require especially careful treatment. EPA[46] states plainly that pesticide labels are legally enforceable and that "the label is the law." The Worker Protection Standard protects agricultural workers and pesticide handlers, including those who mix, load, apply, or repair equipment with pesticide residues. That means the uploaded documents are correct to insist on human-approved spraying, chemical logs, and explicit no-spray zones. A paper that tries to sell unsupervised autonomous pesticide application as a near-term feature would be much weaker. [47]

Workplace safety also matters during service and setup, not only during autonomous operation. OSHA[48] notes that there are no robotics-specific OSHA standards and that many robot incidents occur during programming, maintenance, testing, setup, or adjustment. That is strong support for lockout-minded service procedures, clear maintenance modes, visible status indicators, tool interlocks, and a hardwired E-stop chain independent of high-level software. [49]

Battery safety should be described in standards language, not marketing language. Pack design should align with IEC 62619 for industrial lithium batteries, and commercial rover packs may also target UL 2271 where the product scope fits light electric vehicle energy storage assemblies. The paper does not need to claim immediate certification, but it should show that the architecture is designed with that pathway in mind: BMS supervision, over-current and short-circuit protection, proper charge control, thermal monitoring, enclosure design, and service-safe connectors. UL[50] and IEC both provide a much stronger basis for reviewer confidence than generic battery-safety language. [51]

Maintenance and lifecycle should be treated as first-class design criteria. The uploaded documents are strongest exactly here: steel-tube or otherwise repairable frames, bolted construction, standard bearings and wheels, local parts sourcing, simple wiring, quick-swap tools, and offline-capable workflows. That is not a "nice to have" for smallholder contexts; it is essential. A good paper should therefore specify single-person serviceability for common tasks, tool-less module swaps where feasible, connector standardization, self-test at boot, consumable tracking, and preventive maintenance intervals tied to wheel hours or mission count rather than only calendar time.

The environmental case is favorable but conditional. Precision agriculture can reduce inputs and improve environmental performance; GAO explicitly identifies reduced fertilizer, herbicide, fuel, and water use as a major benefit class. FarmDroid's market case shows that lightweight, precise, renewably powered field robots can substantially reduce chemical use in specific workflows. At the same time, soil work on lightweight robots is not impact-free: repeated wheeling still affects pore functionality and can compact topsoil if poorly managed. A rigorous paper should therefore argue that the rover's environmental benefit depends on mission selection, traffic discipline, and high utilization, not simply on the fact that it is electric. [52]

## Recommendations, conclusion, and open questions

The paper should make five recommendations.

The first is conceptual: position the rover as a bounded field assistant. That wording is stronger than “autonomous farm replacement” because it matches both present technical maturity and the safety standards landscape. [53]

The second is technical: use RTK-GNSS + IMU + wheel odometry + row vision as the baseline localization stack, with SLAM treated as an optional degraded-mode extension. That is simpler, more reproducible, and better aligned with open-field microfarming than a SLAM-first architecture. [54]

The third is product sequencing: launch with seeding, fertilizer micro-dosing, scouting, and mechanical weeding. These modules have the best value-to-risk ratio and the clearest agronomic and regulatory path. Deferred modules should include autonomous spraying beyond human-approved spot work and dexterous harvesting. [55]

The fourth is software governance: treat VLM/VLA technologies as explainability and exception-management layers, not as the sole safety authority for actuation. That preserves the uploaded documents’ explainable-autonomy intent while keeping safety and tool execution in auditable deterministic software. [28]

The fifth is commercial strategy: default to cooperative and service-provider deployment for smallholders, with direct ownership reserved for higher-value and higher-utilization situations such as peri-urban vegetable operations and dense market-garden use. [41]

The conclusion is therefore qualified but clear. Under the stated assumptions, a bounded, semi-autonomous, modular microfarm rover is technically viable and can be economically viable. It is most persuasive when it is lightweight, geofenced, supervised, LiFePO4-electric, solar-assisted rather than solar-dependent, ROS 2-based, RTK-localized, and launched with seeding, micro-dosing, scouting, and weeding modules. That argument is consistent with standards, smallholder mechanization evidence, current robotics software capabilities, field-perception literature, and existing agricultural-robot case studies. The premise is strongest not when it overreaches, but when it is disciplined. [56]

Open questions remain. The most important are the exact first-crop mix, local labor rates, whether the first commercial market is urban/peri-urban or rural smallholder, whether the platform must be chemical-capable in its first release, and what minimum after-sales service network is feasible in the target region. None of those questions undermine the concept. They are the variables that determine which deployment model, module sequence, and capex target should be chosen first.

# Appendices

## Appendix of standards, datasheets, and case sources

| Source family | Specific source | Why it matters |
|---|---|---|
| Agricultural autonomy safety | ISO 18497 parts 1–4 [53] | Machine design, obstacle protection, operating zones, V&V |
| Agricultural control-system safety | ISO 25119 parts 1–4 [57] | Safety-related parts of control systems |
| Workplace robot safety context | OSHA robotics standards and hazard pages [49] | Service, setup, guarding, lockout implications |
| Chemical-use compliance | EPA pesticide labels and WPS materials [47] | Human-supervised spraying logic |
| Battery safety | IEC 62619 and UL 2271 [51] | Pack safety and certification pathway |
| Software stack | ROS 2 docs, Nav2 GPS localization, collision monitor, ROS 1 EOL [58] | Modern open robotics baseline |
| GNSS | u-blox ZED-F9P datasheet, Emlid Reach M2 / RS2+ [59] | RTK/localization baseline |
| Compute | NVIDIA Jetson Orin Nano Super, Raspberry Pi 5 [60] | Edge-compute trade space |
| Vision / sensing | Luxonis OAK-D Lite, Intel RealSense D455, SLAMTEC RPLIDAR A2, Hokuyo UST-10LX [61] | Perception and obstacle-sensing tradeoffs |
| Power | Bioenno 48V 30Ah LiFePO4, Victron MPPT, IEA battery chemistry analysis [62] | Battery and solar-charging design |
| Farm-robot case study | FarmDroid FD20 product and how-it-works pages [33] | Real commercial precedent for solar-powered seeding and weeding |
| Smallholder mechanization | FAO sustainable mechanization and hire-service materials; CGIAR/CIMMYT service-provider model [63] | Business-model and adoption logic |
| Precision-agriculture economics | GAO precision agriculture assessment [34] | Benefits, adoption barriers, data and standards issues |
| Navigation and perception | Agricultural SLAM, GNSS/IMU/vision fusion, crop-row detection [64] | Localization and row guidance evidence |

| Source family | Specific source | Why it matters |
|---|---|---|
| literature | | |
| Task-specific ML literature | Weed/crop detection and edge inference papers; harvesting reviews [65] | Module sequencing and ML selection |
| Soil / environmental literature | Lightweight robot compaction and controlled-traffic studies [9] | Environmental boundary conditions |

## Appendix of modeled economic interpretation

The capex and ROI tables in this report are scenario models built from the cited component price points plus explicit engineering assumptions about fabrication, maintenance, and utilization. They are suitable for a research paper's analytical section, but they should not be presented as observed market averages without a field pilot or customer-discovery dataset.

---

[1] [26] [64] AGRI-SLAM: a real-time stereo visual SLAM for agricultural environment | Autonomous Robots | Springer Nature Link

https://link.springer.com/article/10.1007/s10514-023-10110-y?utm_source=chatgpt.com

[2] [57] ISO 25119-1:2018 - Tractors and machinery for agriculture and forestry — Safety-related parts of control systems — Part 1: General principles for design and development

https://www.iso.org/standard/69025.html?utm_source=chatgpt.com

[3] [8] [29] [46] [56] [63] What FAO does | Sustainable Agricultural Mechanization | Food and Agriculture Organization of the United Nations

https://www.fao.org/sustainable-agricultural-mechanization/overview/what-fao-does/en/?utm_source=chatgpt.com

[4] [20] [32] [45] [53] ISO 18497-1:2024 - Agricultural machinery and tractors — Safety of partially automated, semi-autonomous and autonomous machinery — Part 1: Machine design principles and vocabulary

https://www.iso.org/standard/82684.html?utm_source=chatgpt.com

[5] [12] [14] [58] ROS 2 Documentation — ROS 2 Documentation: Humble documentation

https://docs.ros.org/en/humble/?utm_source=chatgpt.com

[6] [10] [17] [18] [22] [24] [42] [60] Jetson Orin Nano Super Developer Kit | NVIDIA

https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/jetson-orin/nano-super-developer-kit/?utm_source=chatgpt.com

[7] [19] [41] [43] [50] Hire services as a business enterprise: A training manual for small-scale mechanization service providers | Sustainable Agricultural Mechanization | Food and Agriculture Organization of the United Nations

https://www.fao.org/sustainable-agricultural-mechanization/resources/publications/details/en/c/1145335/?utm_source=chatgpt.com

[9] Impacts of load and repeated wheeling from a lightweight autonomous field robot on the physical properties of a loamy sand soil - ScienceDirect

https://www.sciencedirect.com/science/article/pii/S0167198723001587?utm_source=chatgpt.com

[11] [13] [23] [48] [51] IEC 62619:2022 | IEC

https://webstore.iec.ch/en/publication/64073?utm_source=chatgpt.com

[15] [49] Robotics - Standards | Occupational Safety and Health Administration

https://www.osha.gov/robotics/standards?utm_source=chatgpt.com

[16] [34] [52] Precision Agriculture: Benefits and Challenges for Technology Adoption and Use | U.S. GAO

https://www.gao.gov/products/gao-24-105962?utm_source=chatgpt.com

[21] [62] 48V 30Ah LiFePO4 Battery – Compact High-Output Power | BLF-4830A Bioenno Power

https://www.bioennopower.com/products/48v-30ah-lfp-battery-pvc-blf-4830a?utm_source=chatgpt.com

[25] [54] Navigating using GPS Localization — Nav2 1.0.0 documentation

https://docs.nav2.org/tutorials/docs/navigation2_with_gps.html?highlight=gps&utm_source=chatgpt.com

[27] [30] Collision Monitor — Nav2 1.0.0 documentation

https://docs.nav2.org/configuration/packages/configuring-collision-monitor.html?utm_source=chatgpt.com

[28] [44] [2307.15818] RT-2: Vision-Language-Action Models Transfer Web Knowledge to Robotic Control

https://arxiv.org/abs/2307.15818

[31] Electric vehicle batteries – Global EV Outlook 2025 – Analysis - IEA

https://www.iea.org/reports/global-ev-outlook-2025/electric-vehicle-batteries?utm_source=chatgpt.com

[33] FarmDroid FD20 | Autonomous Farming Robot for Precision Agriculture

https://farmdroid.com/products/farmdroid-fd20/?utm_source=chatgpt.com

[35] 9. Technical specifications

https://www.victronenergy.com/media/pg/Manual_SmartSolar_MPPT_75-10_up_to_100-20/en/technical-specifications.html?utm_source=chatgpt.com

[36] [55] Crop detection technologies, mechanical weeding executive parts and working performance of intelligent mechanical weeding: a review - PubMed

https://pubmed.ncbi.nlm.nih.gov/38550283/?utm_source=chatgpt.com

[37] ROS: Upcoming ROS 1 End of Life

https://www.ros.org/blog/noetic-eol/?utm_source=chatgpt.com

[38] [65] Field-based multispecies weed and crop detection using ground robots and advanced YOLO models: A data and model-centric approach - ScienceDirect

https://www.sciencedirect.com/science/article/pii/S2772375524001436?utm_source=chatgpt.com

[39] Sustainable Agricultural Mechanization | Food and Agriculture Organization of the United Nations

https://www.fao.org/sustainable-agricultural-mechanization?utm_source=chatgpt.com

[40] Soil Moisture-Based Irrigation Controllers | US EPA

https://www.epa.gov/watersense/soil-moisture-based-irrigation-controllers?utm_source=chatgpt.com

[47] Introduction to Pesticide Labels | US EPA

https://www.epa.gov/pesticide-labels/introduction-pesticide-labels?utm_source=chatgpt.com

[59] □ ZED-F9P-04B High precision GNSS module Professio

https://content.u-blox.com/sites/default/files/ZED-F9P-04B_DataSheet_UBX-21044850.pdf?hkey=EF798316E3902B6ED9A73243A3159BB0&utm_source=chatgpt.com

[61] OAK-D Lite – Luxonis

https://shop.luxonis.com/products/oak-d-lite-1?utm_source=chatgpt.com