COSMIC

THE POWER OF COLLABORATION

COSMIC Capstone Challenge: Final Briefing

Astrobotics, Purdue University: VLA Robotic Arm Framework

Students: Ankur Senapati, Elaine Lee, Ojas Chaturvedi, Ritwik Jayaraman, Ekansh Agarwal
Advisor: Elizabeth Kung
Mentor: Harsha Malshe

April 13, 2026

ASTROBOTICS

Executive Summary

ASTROBOTICS

Problem:
Current ISAM robotic systems require continuous human-in-the-loop involvement, introducing operational latencies from communication delays, driving up mission costs, and limiting scalability of on-orbit operations.

Proposed Capability:
A platform-agnostic, fully autonomous software framework using foundation vision models (LLaVa) that enables robotic manipulators to execute complex, multi-stage assembly tasks with only high-level natural language supervision.

How It Solves the Problem:
- Eliminates pre-programmed scripts
- Natural language command can be issued by the operator
- The system fuses vision + force sensing to autonomously execute tasks
- Three chained operations (pick-and-place, riveting, screwing) demonstrate a complete assembly capability for operations inside the Bosuns Locker

Status:
- Conceptual design complete (CAD drawings)
- Trade study between VLA and DINO architectures complete
- POC demonstrations of individual operations via simulation complete

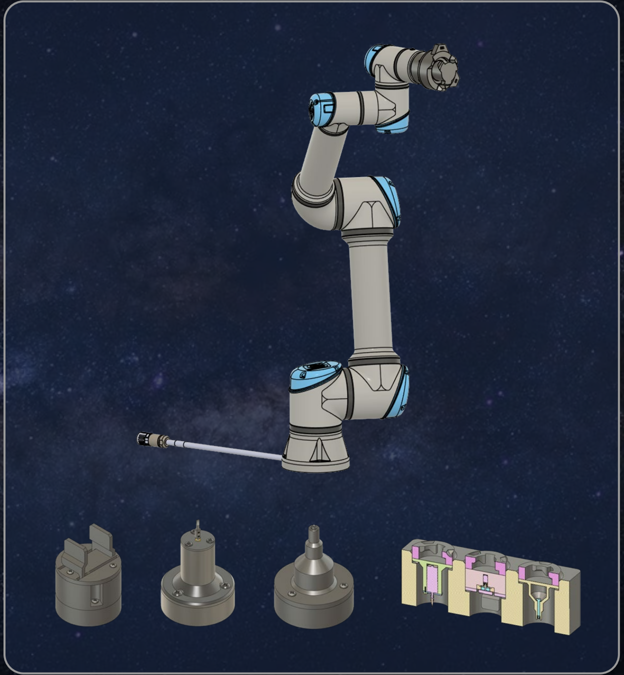

VLA Robot Arm

Team Overview

VLA Robotic Arm Framework

ASTROBOTICS

Ritwik Jayaraman

Elaine Lee

Ojas Chaturvedi

Ankur Senapati

Ekansh Agrawal

2.4 Program Management Milestones

ASTROBOTICS

Category 2: Systems & Program Overview

Sept 2025	Selected Program Manager
Oct 2025	Chose operations: Pick-and-Place, Riveting, Screwing
Oct 2025	Statement of Intent submitted
Nov 2025	Created Systems Requirements Review (SRR)
Dec 2025	Midpoint Showcase pitch delivered
Jan 2026	Completed trade studies (VLA vs DINO, selected VLA/LLaVa)
Mar 2026	Completed simulations & established VLA feasibility
Apr 2026	Finished technical paper & final presentation briefing

The Problem: Why Autonomous Assembly Matters

ASTROBOTICS

$350M+

Cost to replace a single GEO satellite

10-1

Training hours-EVA hour ratio for astronauts

The core challenge:

- Current ISAM system issues:
  - Dependent on continuous human-in-the-loop control
  - Every command must travel from ground to orbit and back, introducing latency that makes real-time precision assembly impractical
- Factors limited by dependency:
  - Mission cadence
  - Increases costs, both financially and time-wise
  - Prevents the scalable orbital manufacturing infrastructure

What is needed:

- A software-first approach:
  - An intelligent autonomy layer that sits between a high-level operator command and the robot's actuators
  - Enables the robot to perceive, plan, and execute multi-step assembly tasks independently with only supervisory oversight

2.2 Storyboard: Concept of Operations

Category 2: Systems & Program Overview

ASTROBOTICS

1. Launch

Payload stowed inside Bosuns Locker, secured in folded config. Launches as hosted payload on Arkisys Port Module.

2. Deploy

Port Module reaches target LEO orbit. Bosuns Locker activated, arm unfolds and performs self-checkout diagnostics.

3. Initialize

Compute unit boots, VLA model (LLaVa) loads. Camera systems verify workspace. System reports ready status to ground.

4. Command

Operator issues high-level natural language task via uplink: "Assemble strut assembly Alpha."

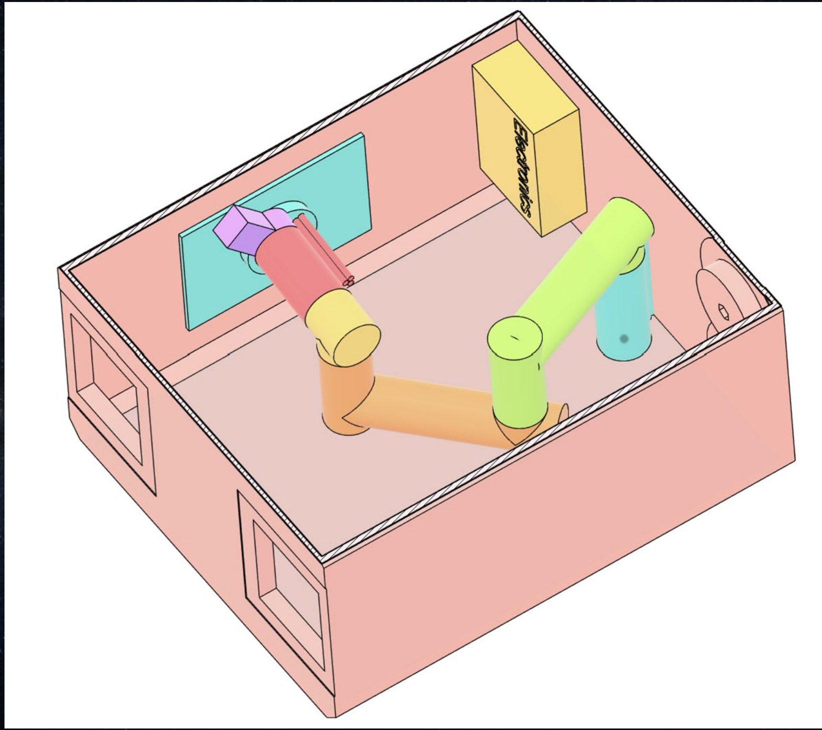

VLA Robot Arm

2.2 Storyboard: Concept of Operations

Category 2: Systems & Program Overview

5. Pick & Place

Arm autonomously identifies component, computes grasp, picks, transports, and places at target location.

6. Rivet

Tool head swaps to rivet attachment. Arm aligns, applies controlled SPR force, verifies joint via force feedback.

7. Screw

Tool head swaps to screw driver. Arm locates fastener points, inserts and torques to spec with closed-loop monitoring.

8. Verify & Report

Vision system inspects completed assembly. Telemetry and images downlinked. Operator confirms or commands next task.

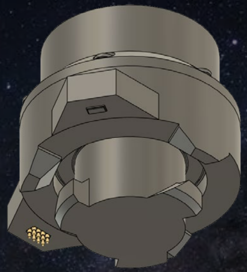

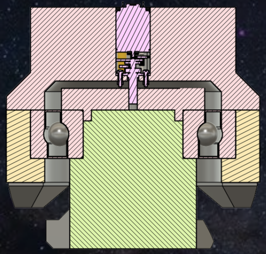

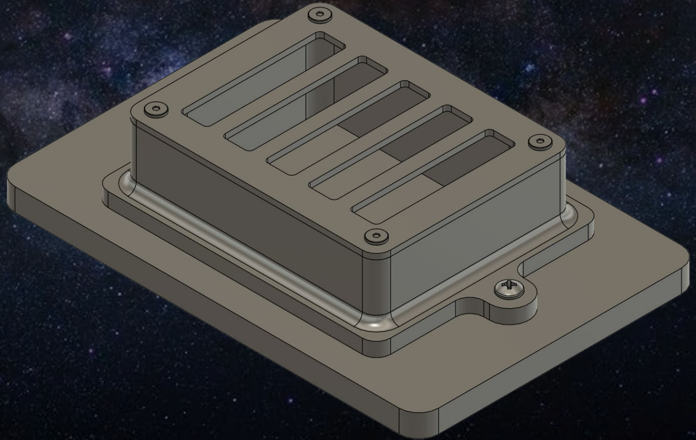

VLA Robot Arm

Three Operations: The Assembly Capability Chain

Our payload demonstrates a chain of 3 discrete operations forming a complete autonomous assembly capability:

01 Pick and Place

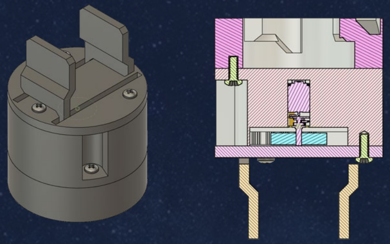

The arm uses camera vision to identify a target component, compute grasp coordinates, approach, grip, transport, and precisely place it at the designated assembly location.

02 Riveting

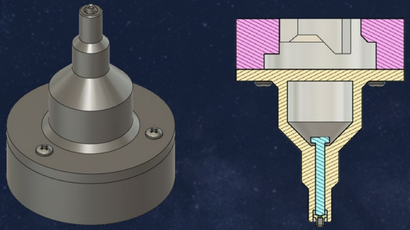

Low-power self-piercing riveting joins structural components without pre-drilled holes. The arm aligns the rivet tool, applies controlled force, and verifies the joint integrity via force feedback.

03 Screwing

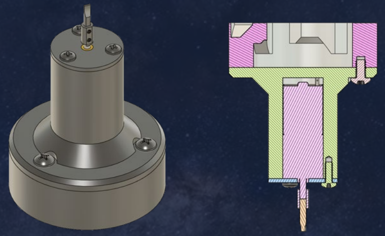

Automated screw-driving using a rotary end effector attachment. The arm locates the fastener point, aligns, inserts, and torques to specification with closed-loop force monitoring.

Technical Overview: The VLA Software Pipeline

Operator issues a natural language command → system autonomously executes the full assembly task

1 VLM Task Reasoning

NL command + camera image → tool selection and task plan

2–4 s inference

2 Visual Detection

Camera frame → Object Detection

Sub-pixel accuracy

3 Depth Deprojection

Pixel coords + depth → 3D world-frame coordinates

< 0.1 cm error

4 Inverse Kinematics

Target end-effector pose → joint angle solution

4 mm accuracy

5 Motion Planning

Current config → collision-free joint-space trajectory

< 5 s planning

6 Smooth Execution

Trajectory → smooth arm motion -> task completed

32 s cycle time

VLA Robot Arm

1.2 Feasibility of Proposed Mission

ASTROBOTICS

Could this mission be designed and built within 5 years?

Bosuns Locker Compliance:

- The 6-DOF arm with end effectors and compute package fits within the 15.75” × 15.75” × 35.45” volume
- Total payload mass estimated well under 400 kg limit
- Sustained power draw for compute + arm actuation within 300W continuous budget.

Technology Readiness:

- Foundation vision models (LLaVas, VLAs) are TRL 4–5 for terrestrial robotics
- Self-piercing riveting is mature (automotive TRL 9)
- The integration challenge is adapting the ML inference pipeline to space-rated compute (e.g., NVIDIA Jetson hardened variants) and validating in microgravity

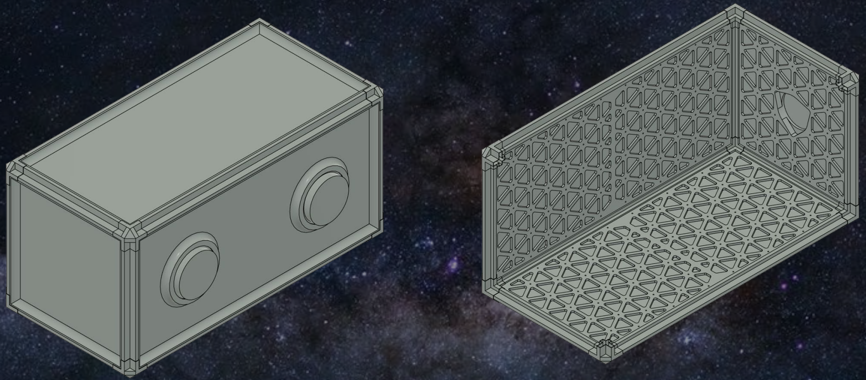

1.2 Feasibility of Proposed Mission

ASTROBOTICS

Could this mission be designed and built within 5 years?

Path to Flight:

- Year 1–2: Ground testbed validation, model optimization for edge compute
- Year 2–3: Thermal/vacuum testing, vibration qualification
- Year 3–4: Integration with Bosuns Locker/Port architecture
- Year 4–5: Flight unit build, integration & test, launch

Bill of Materials:

- 6-DOF robotic arm (UR5 used as a generalized example)
- 3x toolings (Rivet tool head, screwdriver head, claw)
- Structural housing/toolbox
- Software-related components:
    - Cameras
    - On-board computer (including hardened compute, etc.)
    - RAM + long-term memory storage

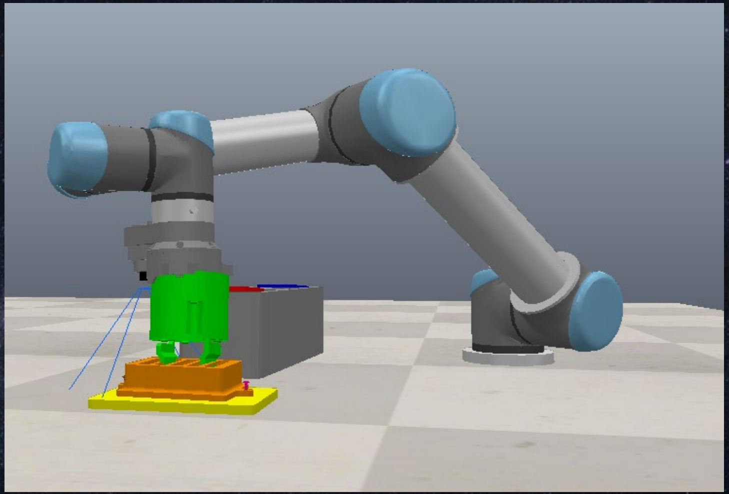

VLA Robot Arm

3.3 Data Handling & Communications

Category 3: Project Engineering

Operational Mode:

- Store-and-forward with periodic real-time telemetry windows
- System operates autonomously after receiving initial task command (does not require a continuous ground link)

Uplink Requirements (Ground → Payload):

- Low bandwidth: Natural language commands (~1 KB per command)
- Supervisory confirmations (proceed/hold)
- No real-time joystick or joint-level control needed

Downlink Requirements (Payload → Ground):

- Moderate bandwidth: Compressed camera images for verification (~500 KB each)
- Telemetry logs (joint positions, force readings, model confidence) ~50 KB/min
- Post-task assembly inspection images
- Estimated total: ~5–10 MB per assembly cycle

Operator Role:

- Supervisory only
- Initiates task sequences
- Reviews post-task verification images
- Authorizes next operation if required

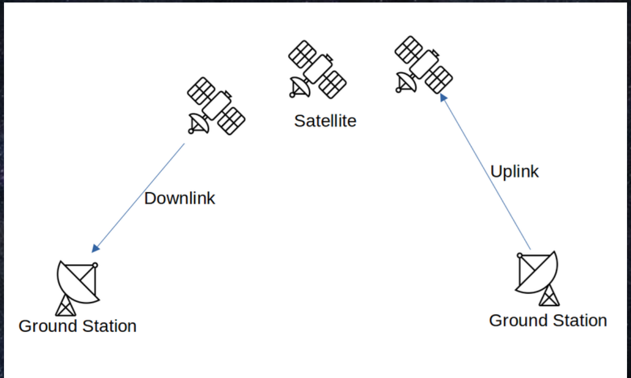

3.1 Completion of Required Elements

ASTROBOTICS

Category 3: Project Engineering

● (Complete)	**Fits within assigned volume?**	6-DOF arm + compute + tool heads package within 15.75" x 15.75" x 35.45" Bosuns Locker envelope. CAD model verified.
● (Complete)	**Within mass & power?**	Estimated payload mass ~80 kg (400 kg limit). Peak draw ~250W; sustained ~180W (300W limit). Power budget includes arm actuators, compute, cameras, end effectors.
● (Complete)	**Bus supports operations?**	Bosuns Locker provides power, data, and structural interface via Port Module. Robotics compatibility built into Arkisys architecture. Pogo pin end effectors designed for Port interfaces.
● (In Progress)	**Launch/descent environment?**	Payload designed with launch restraint brackets. Arm stowed in folded configuration. Vibration analysis for Falcon-class launch loads. Components selected for vacuum and thermal cycling.
● (Complete)	**Demonstrates useful capability?**	Three operations (pick-and-place, riveting, screwing) chain into a complete structural assembly workflow. Each operation leverages the vision model autonomously.
● (In Progress)	**Appropriate analysis?**	Thermal analysis for compute heat dissipation. Structural FEA for arm loads. Software latency analysis for inference pipeline. Power budget analysis.

● Complete ● In Progress

1.1 Technical Impact of Demonstrated Capability

If successfully executed, this mission would:

Enable new types of spacecraft & orbital operations:

- Autonomous multi-step assembly removes the needed human control from orbital construction
- Enables persistent, high-cadence manufacturing operations
- Assembling structures, outfitting payloads, & fabricating components without crew scheduling constraints nor safety concerns

Improve existing spacecraft & ISAM operations:

- Platform-agnostic software can be deployed on any robotic arm
    - Allows existing servicing vehicles to gain autonomous assembly capability through a software upgrade rather than hardware redesign
- Dramatically lowers the barrier to upgrading legacy systems → saves millions of dollars

Reduce dependency on other technology developments:

- Since VLA framework is software-defined and hardware-agnostic, does not require novel manipulator hardware
- Leverages existing robotic arm architectures and COTS compute
- Can be demonstrated near-term without waiting for next-generation space robotics hardware to mature

1.3 Innovation

ASTROBOTICS

First application of foundation vision models to ISAM assembly

No existing on-orbit robotic system uses VLA or LLaVa-class models for autonomous task execution. Current systems rely on pre-scripted motion sequences or teleoperation. Our approach is fundamentally different: the robot reasons about its environment in real time.

Platform-agnostic software architecture

The core intelligence layer is ***decoupled from the hardware***. The same trained model can command different robotic arm configurations with minimal retraining, maximizing reusability across ISAM missions and reducing per-mission software development costs.

Natural language task interface

Operators issue commands ***in plain English*** rather than programming joint trajectories. This collapses the operator expertise barrier — a mission controller does not need to be a robotics programmer to direct complex assembly sequences.

Multi-operation chaining without scripting

The system autonomously sequences pick-and-place → riveting → screwing with ***minimal human in the loop*** between operations. The vision model maintains situational awareness across the full assembly workflow.

1.4 Advancing High-Value Missions

ASTROBOTICS

How does our autonomous assembly capability advance COSMIC's priority ISAM missions?

Autonomous Payload Swap on Persistent Platforms	Our pick-and-place + fastening chain directly enables removing and replacing payloads hosted on persistent platforms like the Arkisys Port. The vision model identifies the payload, the arm extracts it, and installs a replacement, all autonomously with only a high-level command.
Assembly of Large Persistent Platforms	Our software framework enables a robotic arm to autonomously join modular structural components using riveting and screwing, scaling from Bosuns Locker demonstration to full platform assembly
Upgrade/Replacement of Instruments	The same autonomous pick-and-place + fastening capability chain can replace failed or outdated instruments on client spacecraft. Our platform-agnostic software means this capability deploys on any servicing vehicle's robotic arm.
In-Space Assembly of Modular Spacecraft	Assembling spacecraft from modular components in orbit requires precisely the autonomous multi-step manipulation our system demonstrates: identify component, transport, align, join permanently (rivet) or reversibly (screw), and verify

2.1 Animation of Key Operating Sequence

Category 2: Systems & Program Overview

ASTROBOTICS

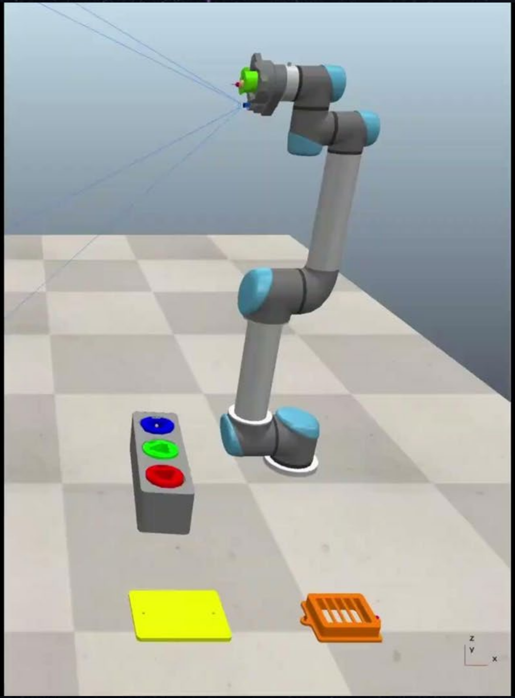

VLA Robot Arm

3.4 Trade Studies

ASTROBOTICS

Trade	Option A	Option B	Decision
Vision Model	VLA: End-to-end but ~10x compute cost. Requires GPU-class hardware.	DINO+Policy Head: Modular, 3x lighter inference. Runs on edge compute.	**VLA/LLaVA selected. End-to-end reasoning outweighs compute cost; optimizable via quantization.**
Joining Method	Welding: Strong permanent joints but extreme heat generation, debris risk in microgravity.	Self-Piercing Riveting: Low power, no pre-drilling, vibration resistant. Proven automotive TRL 9.	**SPR selected. Low thermal signature, no debris, compatible with Bosuns Locker power.**
Reversible Fastening	Adhesives: Simple but not reversible. Outgassing risk in vacuum.	Screwing: Reversible, serviceable, well-understood torque profiles.	**Screwing selected. Enables modular, serviceable structures aligned with ISAM goals.**

3.2 Risk Identification & Mitigation

ASTROBOTICS

Category 3: Project Engineering

Risk	Likelihood	Impact	Mitigation
ML inference latency on space-rated compute	High	High	Target <500ms per inference cycle. Benchmark on Jetson Orin NX or local laptop as proxy. LLaVA architecture supports efficient distillation.
Thermal management of edge compute in vacuum	Med	High	Passive heat sinks with thermal interface to Bosuns Locker structure. Duty-cycle compute to manage thermal budget. Port Module provides thermal services.
Tool changeover reliability in microgravity	Med	Med	Pogo pin magnetic coupling for tool heads. No loose fasteners during swap. Positive locking mechanism with force confirmation sensor
Vision model accuracy degradation from lighting variation	Med	Med	Onboard LED illumination provides consistent lighting. LLaVA vision encoder pretrained on diverse lighting conditions. Fine-tuning on domain-specific data improves robustness.
Communication loss during operation	Low	High	System designed for full autonomy after initial command. Safe-hold state on comm timeout. Stores telemetry for delayed downlink. No real-time ground dependency.

2.3 Path to Preliminary Design Review (PDR)

ASTROBOTICS

Category 2: Systems & Program Overview

Phase 1: Model Validation

- Implement more robust autonomous capability
- Validate on 50+ assembly scenarios in simulation
- Benchmark accuracy vs human teleoperation baseline

Phase 2: Hardware Integration

- Integrate compute unit with flight-representative arm
- Design and manufacture all components (toolings, end-effectors, & attachments)

Phase 3: Testing

- Hardware testing
- Thermal vacuum testing of compute and arm assembly
- Vibration and shock testing
- EMI/EMC testing for Bosuns Locker interface compatibility

Phase 4: System Integration & PDR

- Full system integration with Bosuns Locker mockup
- End-to-end autonomous assembly demonstration
- Complete CDR-level documentation
- Interface verification with Arkisys Port architecture

Estimated timeline from conceptual design to PDR: ~18 months

4.1 Most Innovative Concepts Considered

Category 4: Lessons Learned

1. Sim-to-real transfer for zero-shot orbital deployment

- Train the vision model entirely in simulation (digital twin of Bosuns Locker environment)
- Deploy directly to flight hardware with no on-orbit retraining
- Dramatically reduce commissioning time and eliminate the need for on-orbit data collection before operations begin

2. Multi-arm cooperative assembly via shared foundation model

- VLA models could coordinate two or more robotic arms simultaneously, each arm receiving coordinated action outputs from the same perception backbone
- Enables complex assembly tasks requiring simultaneous holding and fastening

3. Self-supervised anomaly detection during assembly

- Leveraging the VLA feature space to detect assembly anomalies (misaligned parts, incomplete rivets, foreign object debris) without explicit defect training data
- Model would flag deviations from expected visual patterns, enabling quality assurance

4.2 Most Important Technology Gaps

Category 4: Lessons Learned

1. Space-qualified ML inference hardware

- No radiation-hardened processor currently exists that can run transformer-based vision models at the speeds needed for real-time robotic control
- Closing this gap, likely through commercial hardening of edge AI chips (e.g. NVIDIA), would unlock not just our mission but an entire class of AI-driven space robotics

2. Standardized robotic tool-change interfaces for microgravity

- No industry-standard quick-change tool interface exists for microgravity robotic arms
- Standardization of connection ports could benefit ISAM R&D ecosystem

3. On-orbit assembly verification and NDE (Non-Destructive Evaluation)

- No current established method
- Compact NDE sensor (ultrasonic or thermographic) integrated with the end effector greatly increases mission assurance

4.3 Biggest Challenges Encountered

Category 4: Lessons Learned

ASTROBOTICS

1. VLA compute feasibility for space applications (Technical)

- Early benchmarking on powerful embedded hardware revealed inference times of 2–4 seconds per action, providing a need for looking into space hardware solutions
- While DINO was evaluated as a lighter fallback, the superior end-to-end language grounding of the VLA architecture justified the compute investment for complex reasoning

2. Sim-to-real gap in vision model performance (Technical)

- Models trained in simulation initially failed to generalize to physical testbeds due to visual discrepancies in texture, lighting, and environmental reflections
- We utilized domain randomization and fine-tuning on real-world imagery to bridge this performance gap, though it remains a primary risk for flight deployment

3. Cross-disciplinary team coordination (Programmatic)

- Integrating ECE, CS, and Aero team members working on software, controls, and mechanical design respectively required a common “language”
- Initial misalignment in coordinate frames between the CAD model and the vision pipeline cost several weeks of rework

Summary & Conclusion

ASTROBOTICS

What We Built:
A conceptual design for a fully autonomous assembly payload hosted in the Arkisys Bosuns Locker, powered by foundation vision models that eliminate pre-programmed scripts and enable natural language task commanding.

What We Demonstrated:
Three chained autonomous operations — pick-and-place, self-piercing riveting, and screwing — forming a complete structural assembly capability. The system takes a single natural language command and executes the full workflow autonomously.

Why It Matters:
Autonomous software-driven assembly is the key enabler for scaling ISAM from one-off demonstrations to routine orbital manufacturing. Our platform-agnostic approach means this intelligence can be deployed across different robotic systems, accelerating the entire ISAM ecosystem.

Next Steps:
Finetune VLA architecture for different ISAM operations, optimize for edge compute, pursue hardware-in-the-loop testing, and target AIAA SciTech 2027 publication. Explore partnership opportunities to advance toward PDR.

Questions

VLA Robotic Arm Framework

ASTROBOTICS

Questions?

Backup Slides

Backup: VLA vs DINO Architecture Comparison

ASTROBOTICS

VLA (Vision-Language-Action) Model

- Input: Camera image + natural language command
- Output: Direct joint-level action commands
- Architecture: Single transformer-based model (e.g., RT-2, Octo)
- Pros: End-to-end learning, strong generalization across tasks
- Cons: ~2B+ parameters, requires GPU-class compute (~40W+), 2–4s inference on embedded

DINO + Lightweight Policy Head

- Input: Camera image → DINOv2 features + language embedding
- Output: Target end-effector positions (IK solver generates joint commands)
- Architecture: Frozen DINO backbone (~300M params) + small MLP policy (~5M params)
- Pros: 3–5x faster inference, modular (swap policy without retraining vision), lower power (~15W)
- Cons: Requires separate IK solver, less flexible for novel task types

3.3 Data Handling & Communications

Category 3: Project Engineering | 3 points

Ground Station

Operator (supervisory)
Issues NL task commands
Reviews verification images
Authorizes next operation

No joystick / joint-level control

UPLINK
~1 KB / command
NL text only

DOWNLINK
~5–10 MB / cycle
Images ~500 KB ea
Telemetry ~50 KB/min

Bosuns Locker Payload

VLM (LLaVA)
Task reasoning from NL

Vision + Control
OpenCV, IK, OMPL

6-DOF Arm + Tooling
Gripper / Riveter / Screw

Onboard Storage
Telemetry + images

via Arkisys Port Module comms link

Operational Mode: Store-and-forward. Full autonomy after initial command. Safe-hold on comm timeout. No real-time ground dependency. Consistent with C3 "limited remote commands" definition.

Backup: Bosuns Locker Specifications & Compliance

ASTROBOTICS

Parameter	Bosuns Locker Spec	Our Design	Margin
Available Volume	15.75" x 15.75" x 35.45"	~14" x 14" x 32" (stowed)	**Positive**
Payload Mass	400 kg	~80 kg (est.)	**320 kg margin**
Sustained Power	300W	~180W	**120W margin**
Peak Power	1000W	~250W	**750W margin**

All design parameters fall within Bosuns Locker specifications with significant margin.