https://doi.org/10.65166/w703kr86

Frontier Science and Futures
Volume I, Issue 1, July 2026
https://journal.ijhba.com/index.php/fsf
https://slsipress.com

Ariadne's Thread in Deep Space: A Conceptual Framework for AI-Driven, Self-Sustaining, and Communication-Anchored Interstellar Exploration

Ramon George Atento, PhD[1] and Leah Quinto, PhD[2]
[1]First Asia Institute of Technology and the Humanities
[2]De La Salle Medical and Health Science Institute

Abstract

The fundamental incompatibility between human biological limits and the timescales required for interstellar travel renders crewed deep-space mission architectures structurally unreliable across the distances involved (Klee, 2018). Multi-generational ship concepts, while theoretically acknowledged in the literature, introduce compounding failure risks — genetic, psychological, sociological, and mechanical — whose cumulative probability of catastrophic mission failure across centuries of transit approaches certainty (Harkins, 1978; Marin et al., 2022). This paper proposes and analytically evaluates Project Ariadne, a three-pillar conceptual framework for interstellar exploration that eliminates biological survival as a mission-critical dependency by distributing intelligence, physical continuity, and informational persistence across integrated machine systems. The three pillars of the framework are: a fully distributed AI-as-ship architecture in which the spacecraft itself constitutes the mission's operational and adaptive intelligence; a governed swarm-replication system in which population-controlled robotic agents sustain the vessel's physical infrastructure across centuries through layered governance mechanisms including partial replication architecture, apoptosis-analogue decommissioning protocols, and resource-budget constraints; and the Ariadne Thread, a progressive relay-node network deployed along the vessel's trajectory that maintains bi-directional communication between the mission and Earth across interstellar distances. The paper introduces the Checkpoint Horizon as an original theoretical construct, reconceptualizing relay nodes from passive signal repeaters to active mission state archives that preserve recoverable snapshots of mission status at each deployment interval, enabling future Earth-launched missions to reconstruct, locate, and coordinate with the original mission regardless of elapsed time. Drawing on literature in von Neumann probe theory, distributed artificial intelligence, swarm robotics, delay-tolerant networking, and deep-space communication architecture, the paper evaluates each pillar's theoretical grounding, assesses the framework's integrated architectural coherence, and identifies the principal research priorities required for its further development. The paper argues that intelligence, rather than biological presence, is the operative and ultimately sufficient definition of the explorer.

Keywords: *interstellar exploration; distributed artificial intelligence; von Neumann self-replicating probes; governed swarm replication; deep-space relay networks; delay-tolerant networking; autonomous spacecraft architecture*

Article History: Received: March 5, 2026 Revised April 10, 2026 Accepted: May 5, 2026

1. Introduction: The Human Problem in Deep Space

The history of space exploration is, in its deepest architectural sense, a history of life-support engineering. Every crewed mission beyond Earth has required not merely propulsion and navigation, but the continuous biological maintenance of the human beings aboard — their atmosphere, temperature, nutrition, hydration, radiation shielding, psychological stability, and medical care (Childress et al., 2023). For missions within the inner solar system, this engineering challenge is formidable but tractable. For missions to the outer planets, it becomes substantially more demanding (Baratoux et al., 2017; Jones & Kliss, 2009). For missions to the nearest stars, it ceases to be an engineering problem in any conventional sense and becomes a biological impossibility in every practical one (Cardion, 2015).

The nearest stellar system to Earth, Alpha Centauri, lies approximately 4.37 light-years distant (Perryman, 2012). At the velocity of the fastest spacecraft yet launched by humanity — the Parker Solar Probe at its perihelion approach, approximately 0.064 percent of the speed of light (Sebastián et al., 2022) — a vessel would require roughly 6,800 years to complete the journey. Even under the most optimistic near-term propulsion projections, including laser-sail concepts such as Breakthrough Starshot targeting one-fifth the speed of light, transit times to the nearest stellar neighbors range from 20 to 100 years. At more conservative but technologically credible velocities, the figure climbs to several centuries. The biological lifetime of a human being is approximately 80 years under favorable conditions (Rafi & Alavi, 2017). The mismatch is not a matter of degree; it is a matter of kind.

The conventional architectural response to this mismatch has been the multi-generational ship — a closed-system habitat in which successive generations of human beings are born, live, and die aboard a vessel completing a journey that no individual alive at departure will survive to see completed. This solution is not without theoretical merit. It acknowledges the timescale problem honestly and attempts to solve it within a biological framework. However, multi-generational mission architectures introduce compounding failure risks that grow nonlinearly across centuries of operation. Population genetics imposes minimum viable crew sizes, typically estimated between 150 and 500 individuals, to prevent inbreeding depression across 20 or more generations (Smith, 2014; Marin & Beluffi, 2018). Psychological and sociological stability across generations of people who never chose the mission and who have no experiential connection to either origin or destination presents challenges for which no terrestrial analogue exists and for which no validated engineering solution has been proposed. Institutional knowledge degradation, behavioral drift from mission protocols, and the compounding probability of cascading mechanical failures over multi-century timeframes each represent independent axes of catastrophic mission failure. The compound probability that a crewed vessel successfully completes a 500-year transit without irreversible failure — biological, social, mechanical, or informational — is not merely low; it approaches zero.

The problem, stated formally, is this: the biological and temporal fragility of human life is structurally incompatible with the distances involved in deep-space and interstellar exploration. No architecture that places human survival as a necessary condition for mission success can be made reliably robust across interstellar timescales. The only coherent engineering response to this constraint is not to engineer more robust life support, but to remove the life-support requirement entirely — to design missions in which human survival is not a dependency variable.

This paper proposes and develops a conceptual framework that does precisely that. Project Ariadne is a three-pillar theoretical architecture for interstellar exploration in which the mission's intelligence, adaptability, and continuity are distributed across machine systems rather than embodied in human crew. The framework is named for the figure of Greek mythology who provided Theseus with a thread to navigate the labyrinth — not a guide walking beside him, but a persistent connective line, extending back to the point of origin, that both enabled forward progress and preserved the possibility of return, communication, and recovery. The metaphor is structurally apt: Project Ariadne's central engineering intuition is that what a deep-space mission requires is not human presence aboard the vessel, but a persistent, intelligent, and recoverable connection between the vessel and the civilization that launched it.

The three pillars of the framework are as follows. The first pillar is the AI-as-Ship Architecture: a fully distributed artificial intelligence system in which the spacecraft itself constitutes the operational, navigational, and decision-making intelligence of the mission — not a crew managed by onboard computers, but a vessel that is itself the intelligence, organized as a distributed network without a single point of failure. The second pillar is Governed Swarm Replication: a population-controlled model of robotic self-replication aboard the vessel, in which specialized robotic agents capable of manufacturing, repairing, and replacing one another operate under pre-programmed demographic constraints — replication limits, population ceilings, apoptosis-analogue decommissioning protocols — that prevent runaway growth while sustaining the mechanical renewal necessary for centuries-long operation. The third pillar is the Ariadne Thread: a progressive relay-node network in which the vessel periodically deploys semi-autonomous communication anchors at calculated intervals along its trajectory, creating a physical and informational spine that connects the mission to Earth across any distance and across any elapsed time.

These three pillars are mutually reinforcing. The AI-as-ship architecture provides the mission's intelligence and adaptability. The governed swarm-replication system provides the mission's physical continuity and self-repair capacity. The Ariadne Thread provides the mission's informational continuity, its connection to Earth, and the channel through which machine learning updates, navigational corrections, and mission intelligence can be transmitted from an Earth that continues to advance technologically long after the vessel's departure.

The paper further introduces the concept of the Checkpoint Horizon as an original theoretical contribution. Each relay node deposited along the Ariadne Thread functions not only as a passive communication repeater, but as a recoverable informational save point — a complete snapshot of mission status, trajectory, systems health, and accumulated scientific data at the moment of its deployment. The Checkpoint Horizon is the theoretical proposition that any future Earth-launched mission, regardless of how far in the future it departs and regardless of how much faster its propulsion technology may be, can reconstruct the prior mission's full state from any surviving relay node and coordinate with, rendezvous with, or build upon that mission's accumulated progress. The thread, in this sense, does not merely connect — it preserves.

This paper is a theoretical and conceptual contribution; it does not present primary empirical data, experimental results, or engineering specifications. Its method is the structured synthesis and analytical integration of existing literature in four principal domains — von Neumann self-replicating probe theory, distributed artificial intelligence and autonomous spacecraft systems, swarm robotics, and deep-space communication architecture — applied to the evaluation of a coherent original framework. The paper's objectives are as follows: (1) to establish the theoretical foundations, from existing literature, that support each of the three framework pillars individually; (2) to demonstrate the logical and architectural coherence of the three pillars as an integrated mission design; (3) to develop the Checkpoint Horizon concept as an original theoretical construct extending current relay-node and delay-tolerant networking literature; (4) to evaluate the framework against its principal unresolved challenges, including runaway replication risk, interstellar communication bit-rate constraints, and the absence of physical hardware resupply capacity; and (5) to situate the framework within the broader theoretical and philosophical question of what exploration means when the explorer is not human.

The paper proceeds as follows. Section 2 reviews the related literature across five consolidated themes. Section 3 describes the methodology. Section 4 presents the analytical synthesis of the three pillars as a unified architecture and develops the Checkpoint Horizon concept. Section 5 discusses the broader implications of the findings. Section 6 presents conclusions and recommendations for future theoretical and empirical work.

2. Review of Related Literature

2.1 Von Neumann Self-Replicating Probes: Theoretical Foundations, Feasibility Assessments, and Fermi Paradox Implications

The theoretical basis for self-replicating spacecraft derives from John von Neumann's formal work on self-reproducing automata, developed in the late 1940s and published posthumously in 1966 (Cussat-Blanc et al., 2019; Schwartz et al., 1967). Von Neumann demonstrated mathematically that a sufficiently complex machine could, in principle, contain a complete description of itself and use that description to produce functional copies, provided it had access to requisite raw materials and energy (Freitas, 1980). This formalism established that self-replication is not logically paradoxical or physically prohibited — it is, at its core, an information-processing and manufacturing problem (Freitas, 1980). The living cell had already demonstrated as much; von Neumann's contribution was to show that the principle was substrate-independent and therefore applicable, in theory, to machine systems (Langton, 1984).

The first sustained application of this formalism to interstellar exploration was made by Freitas (1980), who proposed a detailed conceptual design for a self-reproducing interstellar probe — subsequently termed a 'von Neumann probe' or 'Bracewell-von Neumann probe' in the literature — capable of traveling to a target stellar system, utilizing local resources to manufacture copies of itself, and dispatching those copies to further stellar destinations. Freitas argued that a civilization deploying even a single such probe could, within geologically modest timescales, achieve complete galactic exploration through exponential replication and diffusion. The logical elegance of this proposal has made it a persistent reference point in both the theoretical spaceflight literature and the broader philosophical discussion of technological civilization capabilities.

Subsequent feasibility assessments have been substantially more cautious. Tipler (1980) accepted the logical validity of Freitas's framework but deployed it as a Fermi Paradox argument: if self-replicating probes are feasible, and if any technological civilization in the galaxy's history has deployed them, their presence should be detectable or their absence requires explanation. Tipler's use of the concept was thus less an engineering endorsement than a logical stress-test. The absence of detected probes was, for Tipler, evidence bearing on the prevalence of technological civilizations — a position that remains contested but has been extensively discussed in the SETI literature (Barlow, 2013; Armstrong & Sandberg, 2013).

More recent assessments have shifted toward near-term engineering appraisal. Borgue and Hein (2021) analyzed the feasibility of near-term self-replicating probes under realistic material and energy constraints, concluding that while full self-replication remains beyond current engineering capability, partial self-replication — in which a probe manufactures specific subsystems or replacement components rather than complete copies — is theoretically achievable within foreseeable technological trajectories. Ellery (2022) advanced this position further, arguing that self-replicating probes are not merely theoretically possible but practically imminent, grounded in the convergence of additive manufacturing, in-situ resource utilization (ISRU), and autonomous robotics research. Ellery explicitly addressed the infrastructure requirements for self-replication aboard a space-based system, identifying power supply, feedstock processing, and computational architecture as the three principal engineering bottlenecks.

Matloff (2022) approached the problem from a propulsion and mission timing perspective, examining how von Neumann probe deployment could be integrated with realistic interstellar propulsion scenarios and how replication timelines at target systems affect the overall galactic diffusion rate of a probe network. Matloff's analysis reinforced the conclusion that even conservative replication rates, combined with modest interstellar velocities, produce galactic-scale coverage within timeframes short relative to the age of the galaxy — a finding with significant implications for both mission architecture and Fermi Paradox reasoning.

Armstrong and Sandberg (2013) provided perhaps the most rigorous modern treatment of the galactic colonization timescale problem, demonstrating through formal modeling that a civilization deploying self-replicating probes at even a small fraction of the speed of light could achieve complete galactic coverage within tens of millions of years — a period orders of magnitude shorter than the galaxy's age. Their work established a firm theoretical upper bound on what self-replicating probe deployment implies for observable galactic signatures and, by extension, what the absence of such signatures implies for civilizational prevalence.

The literature in this theme establishes a coherent theoretical lineage supporting the logical and physical admissibility of self-replicating spacecraft as an interstellar exploration strategy. It simultaneously identifies the principal unresolved constraints — resource acquisition at target locations, the replication bottleneck of complex subsystems, energy requirements for manufacturing operations, and the governance problem of preventing uncontrolled proliferation — that any framework incorporating self-replication must address (Borgue & Hein, 2021; Ellery, 2022). Project Ariadne's engagement with this literature is not to assert that the engineering problems are solved, but to position its governed swarm-replication pillar as a theoretically grounded response to the most tractable version of the self-replication requirement: not galactic-scale probe diffusion, but contained, population-controlled mechanical renewal aboard a single vessel over a centuries-long transit.

2.2 Governed Replication and Population Control in Autonomous Machine Systems

The theoretical elegance of self-replicating systems is matched in the literature by a persistent concern regarding their governance: if a machine system is capable of producing copies of itself, what prevents those copies from replicating without limit, consuming available resources, and ultimately overwhelming any designed operational purpose? This concern — most vividly articulated in Drexler's (1986) 'grey goo' scenario involving molecular nanotechnology — has been a consistent counterweight to enthusiasm for self-replicating machines in both scientific and public discourse. The scientific literature has, however, developed a set of theoretical responses to the uncontrolled replication problem that are substantive, if not yet experimentally validated at the scale required for interstellar applications.

The most biologically grounded approach draws on apoptosis — programmed cell death — as a design template for machine governance. In biological systems, apoptosis functions as a population-regulatory mechanism: cells are genetically programmed to initiate self-destruction after a defined number of divisions, upon detecting specific damage markers, or in response to population-density signals from surrounding cells (Elmore, 2007). The application of this principle to autonomous machine systems proposes that each robotic agent be embedded with a replication counter and a decommissioning protocol triggered when that counter reaches a defined limit, when cumulative operational damage exceeds a threshold, or when population sensors detect that the total agent count has reached a pre-programmed ceiling. This approach does not require centralized oversight; the governance is distributed, embedded in each agent's operational logic, and functions without reference to a supervising controller — a property particularly significant for deep-space applications where communication latency precludes real-time human intervention.

Population dynamics modeling has provided a formal analytical framework for evaluating the stability of such governance mechanisms. The application of Lotka-Volterra predator-prey equations to self-replicating probe populations, as explored in the computational biology literature and applied to SRP mutation dynamics (Langeveld et al., 2022), demonstrates that when replication rates are coupled to resource availability and population density through appropriate feedback functions, the resulting population dynamics tend toward stable equilibria rather than unbounded growth — provided the feedback coupling coefficients are correctly calibrated. This is a significant theoretical result: it suggests that ecological population-regulation principles, which have been extensively validated in biological systems, are in principle transferable to machine population governance, and that stability is achievable without requiring each individual agent to have global knowledge of the system state.

Sagan and Newman (1983) addressed the governance problem from a different angle, examining self-replicating probes in the context of their potential danger to existing ecosystems and civilizations. Their analysis — while framed primarily as a Fermi Paradox argument — articulated the principle that any civilization sophisticated enough to deploy self-replicating probes must also be sophisticated enough to implement effective replication constraints, and that the absence of constraint mechanisms would constitute an engineering failure incompatible with the claimed sophistication of the deploying civilization. This argument, though normative in its framing, establishes an important logical point: governance is not merely desirable but is a design requirement if self-replicating systems are to function as instruments of exploration rather than agents of uncontrolled consumption.

The concept of partial self-replication offers an architecturally distinct approach to the governance problem. Rather than designing agents capable of producing complete copies of themselves — the full von Neumann replicator — partial self-replication constrains agents to manufacturing specific subsystems, components, or replacement parts, with complete agent assembly requiring the coordinated output of multiple specialized agents (Borgue & Hein, 2021; Ellery, 2022). This architecture introduces a structural governance mechanism: no single agent can replicate unilaterally because complete replication requires inter-agent coordination that is itself subject to population-level oversight protocols.

Resource-limit governance represents a third distinct mechanism. In this model, replication is constrained not by agent-internal programming but by the finite availability of the materials required for agent manufacture. If the vessel's raw material inventory is managed by a resource allocation system that enforces strict budgetary limits on replication-directed material flows, unbounded population growth is physically impossible regardless of individual agent replication behavior (Ellery, 2022). This mechanism is robust to programming failures — a replication counter error cannot produce runaway growth if the material budget has been physically exhausted — but introduces the complementary risk of material misallocation leading to population decline below mission-critical levels.

The literature on governed replication thus presents a portfolio of theoretical mechanisms — apoptosis analogues, population-density feedback, partial replication architectures, and resource-limit governance — each addressing the uncontrolled growth problem from a distinct engineering angle and possessing distinct failure modes. No single mechanism has been experimentally demonstrated at the scale or in the environmental conditions relevant to interstellar missions. The theoretical basis for governance is, nonetheless, substantially more developed than public discourse around self-replicating machines typically acknowledges, and the convergence of multiple independent mechanisms within a single governance architecture — as Project Ariadne proposes — provides a theoretical redundancy that significantly reduces the probability of governance failure.

2.3 Distributed Artificial Intelligence and Autonomous Spacecraft Architecture

The conceptual foundation for Project Ariadne's first pillar — the proposition that a spacecraft can itself constitute the mission's intelligence rather than carrying intelligence as a separate system — rests on two intersecting bodies of literature: the theoretical and engineering literature on distributed artificial intelligence, and the applied literature on autonomous spacecraft systems and their capacity for unsupervised long-duration operation.

Distributed artificial intelligence (DAI) denotes architectures in which cognitive and computational functions — perception, planning, decision-making, learning, and communication — are distributed across a network of interacting agents or processing nodes rather than concentrated in a single computational system (Weiss, 1999). The defining properties of well-designed DAI architectures are precisely the properties required by a deep-space mission operating beyond the practical range of real-time human oversight: fault tolerance through redundancy, adaptive responsiveness to environmental changes not anticipated at design time, scalability across varying computational loads, and the capacity to maintain coherent system-level behavior from the local interactions of individually limited agents. DAI has been extensively studied in the context of large-scale distributed computing, multi-robot coordination, and networked autonomous systems, and its theoretical properties are well-established (Lesser, 1999; Dorigo & Birattari, 2007).

The application of DAI principles to spacecraft architecture has been most systematically developed through NASA's Autonomous NanoTechnology Swarm (ANTS) program, which proposed a mission architecture for asteroid belt exploration in which a large population of small autonomous spacecraft operated without centralized control, collectively performing distributed sensing, mapping, and data relay tasks through local inter-agent communication (Curtis et al., 2003). The ANTS program is significant for the present framework not as a direct technological precursor but as a proof-of-concept demonstration that mission planners have seriously considered — and developed architectural specifications for — spacecraft systems in which intelligence is a distributed property of the collective rather than a localized function of a single onboard computer.

The broader autonomous spacecraft literature has developed substantially in the two decades since ANTS. NASA's Jet Propulsion Laboratory has pursued swarm autonomy research encompassing decentralized guidance, navigation, and control (GNC) for multi-spacecraft formations, including the development of algorithms enabling spacecraft to autonomously negotiate formation geometry, collision avoidance, and task allocation without ground intervention (Scharf et al., 2004). The Mars Science Laboratory's autonomous driving capabilities (Verma et al., 2023), the AEGIS autonomous targeting system aboard the Mars Exploration Rovers (Estlin et al., 2012), and the experimental onboard science autonomy system of the Earth Observing-1 satellite (Chien et al., 2004) each represent incremental extensions of spacecraft autonomy toward the unsupervised long-duration operational profiles required for interstellar missions.

The question of remote software upgradeability — the capacity to transmit functional improvements to an autonomous system operating at interstellar distances — represents a critical and underexplored dimension of the distributed AI literature in the spaceflight context. Research on federated learning architectures — in which model updates are computed at a central node and transmitted to distributed agents for local integration — provides a theoretical template for packaging and transmitting machine learning improvements for distributed AI systems (Konecny et al., 2016). The communication latency inherent in interstellar relay transmission means that updates would function asynchronously — improving the vessel's capabilities for future decision-making rather than correcting current behavior in real time.

A significant gap in the existing autonomous spacecraft literature, noted by Hein and Baxter in their review of AI in self-replicating probe contexts, is the absence of validated architectural frameworks for AI systems designed to operate autonomously for durations exceeding decades without human maintenance or reprogramming (Hein & Baxter, 2023). Current autonomous spacecraft systems are designed for mission durations ranging from months to a few years; their autonomous operation protocols are developed with the assumption of periodic ground contact, software patching, and mission replanning from Earth (Nesnas et al., 2021). The extrapolation from these systems to a centuries-long fully autonomous mission involves not merely an increase in operational duration but a qualitative shift in architectural requirements — the system must be capable of self-directed adaptation to failure modes, environmental conditions, and mission contingencies that cannot be anticipated at design time and cannot be resolved by ground intervention (Steinberg et al., 2016).

2.4 Swarm Robotics: Decentralized Coordination, Fault Tolerance, and Space Exploration Applications

Swarm robotics is a subfield of multi-robot systems research concerned with the design and analysis of large populations of relatively simple robotic agents whose collective behavior produces complex, adaptive, and robust task execution through purely local interactions, without centralized control or global knowledge (Dorigo et al., 2014). The field draws theoretical inspiration from biological swarm systems — ant colonies, bee swarms, bird flocking, fish schooling — in which individually limited agents collectively achieve sophisticated outcomes through simple local rules and stigmergic communication (Roy et al., 2014). The engineering significance of this approach lies in its scalability and fault tolerance: because no individual agent is irreplaceable and no central controller exists whose failure would disable the system, swarm systems are architecturally robust to the loss of individual agents in ways that centrally controlled systems are not (Brambilla et al., 2013).

The fault tolerance properties of swarm robotic systems have been studied extensively and are among the most consistently supported findings in the literature. Brambilla et al. (2013), in a comprehensive review of swarm robotics research, identify redundancy-based fault tolerance as an emergent property of swarm architectures: because the same functional task is performed by multiple agents operating in parallel, the failure of any individual agent reduces the swarm's performance marginally rather than catastrophically. This property scales favorably with swarm size — larger swarms are more fault-tolerant — and is preserved under conditions of simultaneous multiple-agent failure, provided failures do not exceed a threshold fraction of the total population. The threshold fraction above which swarm coherence degrades is task- and architecture-dependent, but empirical studies across multiple swarm platforms consistently identify substantial robustness margins, with coherent collective behavior maintained even with significant individual agent failure rates, depending on the specific task and communication topology (Devi et al., 2024; Duarte et al., 2016; Dorigo et al., 2014).

Decentralized coordination in swarm systems operates through a combination of direct agent-to-agent communication and stigmergy — the indirect communication mediated by modifications to a shared environment. In robotic swarm systems, stigmergic communication is typically implemented through shared data structures, environmental markers, or physical modifications to a workspace that convey information to subsequently acting agents without requiring direct inter-agent messaging (Dias et al., 2023; Zambonelli et al., 2014; Dorigo & Birattari, 2007). The relevance of stigmergy to the Ariadne framework is significant: aboard a deep-space vessel, the physical state of the vessel itself — the distribution of manufactured components, the status of repair operations, the inventory of raw materials — functions as a persistent stigmergic medium through which robotic agents coordinate their manufacturing, repair, and replication activities.

Task allocation in swarm systems has been studied through threshold-based response models, in which individual agents initiate a task when a perceived stimulus exceeds an agent-specific response threshold (Bonabeau et al., 1999). This mechanism produces adaptive, self-organized task allocation across the swarm without any agent requiring global knowledge of the system's state, and has been demonstrated to produce efficient allocation in dynamic environments where task demands change unpredictably over time (Low et al., 2006).

The application of swarm robotics principles to space exploration has been proposed across several contexts. The NASA ANTS program drew explicitly on swarm principles for distributed spacecraft operation (Hinchey et al., 2007). The Marsbee project, a NASA-funded NIAC study, proposed swarm robotics for atmospheric sensing on Mars, exploiting swarm redundancy to compensate for the high individual failure rates expected in the Martian atmospheric environment (Timmons et al., 2021; Chang et al., 2018). A 2024 review identified the emerging convergence of swarm robotics with in-situ resource utilization as a near-term research priority for lunar and planetary construction applications — a convergence directly relevant to the manufacturing and repair tasks envisioned in the Ariadne framework (Azami et al., 2024; Tan et al., 2024; Chen et al., 2024).

The primary limitations identified in the swarm robotics literature that bear on the Ariadne framework concern the complexity scaling of swarm behavior in highly constrained environments and the challenge of designing reliable emergent behaviors in systems where individual agents are subject to manufacturing variability and cumulative operational degradation. These are recognized limitations in the literature, not disqualifying objections, and they define important research directions rather than theoretical prohibitions.

2.5 Deep-Space Communication Architectures, Delay-Tolerant Networking, and Asynchronous Mission Coordination

The communication infrastructure required to sustain a persistent, bi-directional information link between an interstellar vessel and Earth constitutes one of the most technically demanding components of any deep-space mission architecture. Contemporary deep-space communication is managed through NASA's Deep Space Network (DSN), a ground-based antenna infrastructure capable of maintaining radio-frequency communication links with spacecraft at distances up to the outer solar system. Radio-frequency transmission at interstellar distances produces signal attenuation that renders communication at useful data rates effectively impossible without either impractically large transmission power budgets or receiving antenna arrays of dimensions that are not near-term achievable (Hippke, 2021). Laser communication offers a substantially more favorable ratio of transmitted power to received signal strength due to the narrow beam divergence of coherent optical transmission, and the literature increasingly regards laser communication as the enabling technology for interstellar data links (Marcy et al., 2022).

Hippke (2021) provides the most comprehensive recent treatment of interstellar communication network topology, addressing the optimization of relay node placement across interstellar distances for maximum network throughput, signal preservation, and redundancy. Hippke's analysis demonstrates that a network of relay nodes placed at intervals calibrated to preserve laser beam coherence across each relay-to-relay segment can theoretically sustain meaningful data rates across light-year-scale distances, with each relay node amplifying, error-correcting, and retransmitting the signal to the next. Hippke identifies power and mass budget as the critical engineering constraints governing relay node amplification and retransmission capacity.

Gravitational lensing relay concepts represent a theoretically elegant supplementary communication architecture. Maccone (2009) demonstrated that the Sun's gravitational field produces a focal point at approximately 550 astronomical units (AU) from the Sun, at which the Sun acts as a gravitational lens capable of amplifying received signals by a factor of approximately 100 decibels — an amplification that would render interstellar communication at meaningful data rates achievable with transmitters of practical power levels.

Delay-Tolerant Networking (DTN) addresses the protocol-level challenge of maintaining reliable data delivery over communication links characterized by extreme latency, intermittent connectivity, and high error rates (Cerf et al., 2007). The DTN Bundle Protocol (BP), standardized by the Internet Engineering Task Force and adopted by NASA as the interoperability standard for deep-space networking, implements a store-and-forward architecture in which data bundles are transmitted, stored at each network node until a forward link becomes available, and progressively relayed toward the destination without requiring an end-to-end path to exist at any single moment (Scott & Burleigh, 2007). This architecture is fundamentally compatible with the relay-node topology of the Ariadne Thread: each deployed relay node functions as a DTN store-and-forward node, buffering transmissions from the vessel and from Earth and relaying them forward or backward through the network as link conditions permit.

The scalability of DTN protocols to interstellar distances and multi-decade latencies has been examined theoretically but not empirically validated at such scales (Fraire et al., 2018). Existing DTN deployments operate at communication latencies measured in minutes to tens of minutes (Araniti et al., 2015); the extension to latencies measured in years introduces architectural requirements for persistent storage across timescales that exceed the operational lifetime of current storage media and for error-correction protocols robust to the gradual degradation of relay node electronics over decades of autonomous operation (Caini et al., 2011).

Freitas (1980) briefly addressed the coordination problem between reproductive and non-reproductive probe strategies, noting that the informational infrastructure established by a relay network would enable later missions to benefit from earlier missions' data even without physical rendezvous. The general interstellar mission design literature has addressed multi-mission coordination in the context of mission architecture sequencing (Chen et al., 2021), but has not developed the relay-infrastructure-dependent coordination model that the Ariadne framework proposes as the Checkpoint Horizon concept. This gap represents both a limitation in the existing literature and the primary conceptual space into which the present paper's original theoretical contribution is directed.

2.6 Synthesis of Literature

The five literature themes reviewed above collectively establish a theoretical landscape that is neither uniformly supportive nor uniformly prohibitive of the Project Ariadne framework. Rather, the literature presents a pattern of convergent theoretical feasibility at the level of individual pillars, combined with acknowledged engineering gaps at the level of integrated, long-duration deployment. Several cross-cutting observations emerge from this synthesis.

First, there is a consistent pattern across all five themes of theoretical principles being well-established while operational validation at the required scale is absent. Von Neumann self-replication is logically admissible and has been realized in laboratory settings (Pesavento, 1995); governed replication mechanisms are theoretically coherent but have yet to be demonstrated in autonomous systems of relevant complexity (Sadownik et al., 2017); distributed AI for spacecraft is architecturally plausible but validated only for mission durations orders of magnitude shorter than interstellar transits (Ellery, 2025; Hein & Baxter, 2018); swarm robotics fault tolerance is empirically robust in laboratory conditions but untested in deep-space manufacturing environments (Bossens & Tarapore, 2020; Winfield & Nembrini, 2006); and DTN protocols are operationally proven at solar-system scales but theoretically extrapolated for interstellar latencies (Alhilal et al., 2019; Fraire et al., 2018). This convergent pattern is significant: it means the framework is not confronting isolated speculative claims but rather a consistent frontier condition in which foundational science is settled and engineering extrapolation is the operative challenge.

Second, the literature reveals a structural complementarity among the three pillars that existing reviews have not explicitly identified. The fault tolerance properties of swarm robotics address precisely the individual-agent failure modes that threaten the governed replication system; the distributed AI architecture provides the coordination substrate that swarm robotics requires for complex adaptive task execution; and the DTN relay network provides the update channel through which the AI architecture can receive improvements that compensate for its inability to anticipate all contingencies at design time. The three pillars are not merely co-present in a single mission — they are mutually compensatory, each addressing a principal vulnerability of the others.

Third, the literature's treatment of the Fermi Paradox implications of self-replicating probe deployment establishes an important contextual point: the logical argument that a civilization's failure to deploy self-replicating probes constitutes an inexplicable omission (Ellery, 2022) implicitly positions the governed, bounded version of such deployment — as proposed in the Ariadne framework — as a more defensible and responsible instantiation of the concept than the unconstrained galactic diffusion model that generates Fermi Paradox anxieties.

2.7 Gaps in the Literature

The review identifies four principal gaps that directly motivate the present paper's contribution. First, no existing study has proposed an integrated mission architecture that combines distributed AI, governed swarm replication, and progressive relay-node communication within a single coherent design framework; the three bodies of literature have developed largely independently, with limited cross-referencing. Second, the autonomous spacecraft literature has not addressed AI architectural requirements for mission durations exceeding decades, leaving a critical gap between current autonomous systems design and the centuries-long operational profile of an interstellar mission. Third, the deep-space communication literature has not developed a formal theoretical model of relay nodes functioning as recoverable mission state archives — the informational function that the Checkpoint Horizon concept addresses — treating relay nodes exclusively as signal transmission waypoints rather than as persistent informational repositories. Fourth, the interstellar mission design literature has not systematically examined the coordination mechanics between temporally separated missions connected by a shared relay infrastructure, leaving the asynchronous multi-generational coordination proposition without a formal analytical framework.

2.8 Contribution of the Present Paper

The present paper contributes to the literature in three distinct respects. First, it proposes and develops the first explicit integrated framework combining distributed AI architecture, governed swarm replication, and progressive relay-node communication as mutually reinforcing pillars of a unified interstellar mission design — a synthesis that, to the authors' knowledge, has not been previously formalized in the scholarly literature. Second, it introduces the Checkpoint Horizon as an original theoretical construct, reconceptualizing relay nodes from passive signal repeaters to active mission state archives whose informational function persists independently of the vessel's operational status and enables both remote mission reconstruction and asynchronous coordination with future missions. Third, it provides a structured analytical evaluation of the three-pillar framework against its principal unresolved challenges — governed replication stability, interstellar relay communication constraints, and hardware resupply limitations — establishing the framework's theoretical defensibility while precisely delineating the research agenda required for its further development.

3. Methodology

3.1 Research Design

This paper is a theoretical and conceptual contribution employing a structured thematic literature review as its primary methodological instrument. The design is appropriate to the paper's objectives, which are to synthesize, evaluate, and integrate existing theoretical and empirical literature across four disciplinary domains — self-replicating systems theory, distributed artificial intelligence and autonomous spacecraft architecture, swarm robotics, and deep-space communication engineering — in service of developing and assessing an original conceptual framework. No primary data were collected. No experiments, surveys, interviews, simulations, or statistical procedures were conducted or are reported. The paper does not claim empirical findings; its analytical outputs are theoretical propositions, conceptual integrations, and structured evaluative arguments grounded in the reviewed literature.

The choice of a theoretical conceptual paper as the appropriate design is consistent with the paper's subject matter. Project Ariadne addresses a problem domain — interstellar mission architecture combining self-replicating robotic systems, distributed AI, and relay-node communication networks — in which the absence of operational precedent makes empirical study impossible and the primary intellectual work involves the construction and evaluation of coherent theoretical frameworks from existing foundational literature. This is a well-established methodological mode in frontier science and technology studies, where conceptual framework development frequently precedes and enables the empirical research programs that follow (Andersson et al., 2023; Simões et al., 2022).

3.2 Source Orientation and Inclusion Logic

Source selection was oriented toward peer-reviewed and formally published literature in the natural sciences, engineering, and applied computer science, supplemented by technically rigorous non-peer-reviewed sources — including NASA technical reports, formally published conference proceedings, and established institutional research documentation — where peer-reviewed coverage of a specific technical area was limited due to the frontier character of the subject matter. Given the speculative and forward-looking nature of interstellar mission architecture as a field, certain foundational contributions exist in the form of technical monographs, formally documented research programs, and book-length theoretical treatments rather than journal articles; these were included on the basis of their evidential and theoretical significance rather than excluded on the basis of publication format.

Inclusion criteria for sources were as follows: relevance to one or more of the five literature themes identified in the framework phase; technical or scientific rigor as assessed by publication venue, citation record, institutional provenance, or the methodological transparency of the work's own claims; and sufficient specificity to contribute substantively to the evaluation of the Ariadne framework's three pillars rather than merely providing general background orientation. Sources making extraordinary empirical claims without methodological support, sources confined to popular science presentation without underlying technical content, and sources representing pure speculation without grounding in established scientific principles were excluded.

No formal bibliometric screening protocol — such as PRISMA or SPIDER — was applied, as the structured thematic review design employed here is oriented toward theoretical synthesis rather than systematic evidence aggregation. The literature review is selective and integrative rather than exhaustive, and does not claim to represent a complete census of publications in any of the constituent fields.

3.3 Thematic Grouping Logic

The literature was organized into five themes — von Neumann self-replicating probes; governed replication and population control; distributed AI and autonomous spacecraft architecture; swarm robotics; and deep-space communication architecture and delay-tolerant networking — on the basis of functional correspondence to the three pillars of the Ariadne framework. Themes 1 and 2 together ground Pillar 2 (governed swarm replication), with Theme 1 establishing the theoretical lineage and feasibility envelope of self-replication and Theme 2 addressing the specific governance mechanisms required to make replication operationally viable within the framework. Theme 3 grounds Pillar 1 (AI-as-ship architecture) and its remote upgradeability requirement. Theme 4 provides the operational and coordination substrate shared between Pillar 1 and Pillar 2. Theme 5 grounds Pillar 3 (the Ariadne Thread relay network) and the Checkpoint Horizon concept.

3.4 Evaluative and Analytical Criteria

The analytical synthesis section evaluates the three-pillar framework against three categories of criteria derived from the literature review. The first category is theoretical grounding: the extent to which each pillar rests on established and internally consistent scientific and engineering principles rather than unsupported extrapolation. The second category is architectural coherence: the extent to which the three pillars function as a mutually reinforcing integrated design rather than as independently proposed components that merely co-inhabit a single paper. The third category is challenge tractability: the extent to which the principal unresolved problems identified in the literature represent engineering frontiers with plausible solution pathways rather than theoretical prohibitions.

The paper further distinguishes consistently among four epistemic registers in its analytical claims: established — claims directly supported by experimental or observational evidence in the reviewed literature; theoretically grounded — claims supported by internally consistent theoretical argument and established physical or mathematical principles, but not yet experimentally validated at the relevant scale; extrapolated — claims derived by principled extension from established results in analogous systems or conditions, carrying greater inferential uncertainty; and speculative — claims that are logically coherent and not contradicted by established science but that rest on assumptions sufficiently distant from validated knowledge to require explicit flagging.

3.5 Limitations of the Approach

Several limitations of the present methodology require acknowledgment. First, the absence of empirical data collection means that the paper's evaluative conclusions cannot be validated against operational evidence; the framework's three pillars are assessed for theoretical coherence and literature support, not for demonstrated performance. Second, the structured thematic review design is inherently selective, and the thematic boundaries established here inevitably foreground certain bodies of literature while placing others — including the materials science literature on in-situ resource utilization, the long-duration power systems literature relevant to relay node energy supply, and the astrodynamics literature on trajectory optimization for relay deployment — at the periphery of the review. Third, the frontier character of the subject matter means that a substantial portion of the most technically detailed and recent work exists in preprint, technical report, or conference proceeding form, introducing variability in the peer-review status of cited sources. Fourth, the Checkpoint Horizon concept and the asynchronous multi-generational coordination model are introduced as original theoretical constructs and have not been subjected to independent external evaluation; their theoretical coherence is argued within the present paper but not confirmed by prior literature.

4. Analytical Synthesis: The Ariadne Framework as an Integrated Mission Architecture

4.1 The Three-Pillar Architecture as a Unified System

The analytical work of this section is not the sequential evaluation of three independent proposals but the demonstration that the three pillars of Project Ariadne constitute a structurally integrated system in which each pillar's operational viability depends upon and reinforces the others. This interdependence is not incidental to the framework's design; it is its principal architectural feature and the primary respect in which it differs from prior treatments of distributed AI, swarm replication, and relay communication as separate research problems (Communication Papers of the 17th Conference on Computer Science and Intelligence Systems, 2022; Naylor et al., 2014; Varadharajan et al., 2020).

The AI-as-ship architecture provides the mission's adaptive intelligence — the capacity to perceive the vessel's operational state, interpret environmental data, make navigational and resource allocation decisions, and respond to failure events without human intervention. However, a distributed AI operating across a vessel for centuries without physical maintenance will inevitably confront hardware degradation that its software intelligence alone cannot address (Hein et al., 2020). Component failures, radiation-induced processing errors, and the cumulative mechanical deterioration of sensors, actuators, and computing substrate each represent modes of operational decline that no amount of algorithmic sophistication can compensate for if the physical substrate of the intelligence is irreversibly compromised (Lange et al., 2024; Trindade, 2021). The AI architecture is therefore dependent on the swarm-replication system to maintain the physical integrity of the distributed computing and sensing hardware through which its intelligence is instantiated.

The governed swarm-replication system provides precisely this physical continuity — the capacity to manufacture replacement components, repair damaged subsystems, decommission degraded agents, and sustain the vessel's mechanical and electronic infrastructure across timescales that exceed the operational lifetime of any individual component. However, the swarm's coordination capacity depends on its access to a coherent decision-making framework that can allocate manufacturing priorities, manage resource budgets, and arbitrate between competing repair and replication demands across the vessel's distributed task environment. Without the AI architecture's coordination substrate, the swarm's governed behavior would be confined to the pre-programmed behavioral rules embedded at launch — rules that cannot anticipate the specific failure sequences and resource conditions that will emerge across centuries of operation.

The Ariadne Thread relay network provides the mission's informational continuity — the persistent connection to Earth that enables the transmission of machine learning updates, navigational corrections, and mission intelligence through the relay network to the vessel. However, the relay network's informational value is realized only if the vessel receiving updates is capable of integrating them into its operational intelligence in a functionally meaningful way — a requirement that presupposes the AI architecture's evolvable software design. Conversely, the AI architecture's capacity to remain current with Earth's advancing knowledge depends entirely on the relay network's persistence; without the Thread, the vessel's intelligence is frozen at its launch-time configuration regardless of how many decades of AI research Earth subsequently produces.

This triangular interdependence — AI requiring swarm for physical continuity, swarm requiring AI for adaptive coordination, both requiring the relay network for external intelligence — means that the failure of any single pillar does not merely reduce the mission's performance in one dimension but initiates a cascade of functional degradation across the other two. Conversely, the successful integration of all three pillars produces a mission architecture whose robustness exceeds what any individual pillar could sustain independently, because the system possesses multiple independent recovery pathways for each category of failure.

4.2 Evaluating the AI-as-Ship Pillar

Considering the theoretical grounding criterion established in the methodology, the AI-as-ship pillar is supported by well-established principles of distributed computing and fault-tolerant system design (Fayyaz & Vladimirova, 2016), operationalized in space exploration contexts through the ANTS program and the progressive autonomy extensions demonstrated across NASA's planetary mission portfolio (Truszkowski et al., 2006). The pillar's theoretical grounding is strongest at the level of architectural principle — that distributing intelligence across a network of interacting nodes produces fault tolerance, adaptability, and scalability — and weakest at the level of mission duration, where the extrapolation from mission autonomy validated across years to the centuries-long operational profile required for interstellar transit involves qualitative differences in system requirements that the existing literature does not resolve.

The remote upgradeability sub-component of this pillar is theoretically grounded in the federated learning literature (Lo et al., 2020) and in established principles of modular software architecture (Arafeh et al., 2023), but rests on the additional assumption that the vessel's AI substrate will remain sufficiently intact despite the communication latency of relay-mediated transmission — potentially years between transmission and receipt — for received updates to be coherently integrated. This assumption is reasonable if the swarm-replication pillar maintains hardware integrity, but introduces a dependency that renders the upgradeability proposition contingent rather than unconditional. The epistemic register of the remote upgradeability claim is therefore extrapolated rather than established — a principled inference from analogous systems that requires validation in architectures specifically designed for the interstellar operational context.

4.3 Evaluating the Governed Swarm-Replication Pillar

The governed swarm-replication pillar presents the most complex evaluative profile of the three, because it rests on theoretical mechanisms from two distinct literatures — swarm robotics and governed replication — whose integration produces properties not individually present in either. The swarm robotics literature establishes fault tolerance (Buchanan et al., 2023) and adaptive task allocation (Balachandran et al., 2024) as robust properties of decentralized multi-agent systems, validated empirically across multiple hardware platforms and task environments. The governed replication literature establishes the theoretical coherence of apoptosis analogues, population-density feedback, partial replication architectures, and resource-limit governance as mechanisms for constraining autonomous replication (Ellery, 2022; McMullin, 2000). Each body of literature is, within its own domain, theoretically well-grounded.

The integration challenge is that governed self-replication in a swarm context — where the agents performing replication are themselves subject to swarm coordination dynamics, where manufactured agents may exhibit manufacturing variability relative to their predecessors, and where governance mechanisms must remain effective across populations whose individual members are the products of iterated replication rather than controlled manufacturing — introduces complexity that neither of these literatures has directly addressed. In particular, the degradation of governance fidelity across replication generations — the possibility that apoptosis-analogue programming or population-ceiling constraints may be imperfectly reproduced in daughter agents — represents a failure mode with no direct analogue in either the swarm robotics or governed replication literature as currently constituted.

The paper's evaluative judgment, grounded in the reviewed literature, is that this integration challenge represents an engineering frontier rather than a theoretical prohibition. The most defensible near-term instantiation of the swarm-replication pillar is the partial replication model identified by Borgue and Hein (2021) and Ellery (2022), in which agents manufacture subsystems rather than complete copies, because this architecture structurally limits the governance fidelity problem: no agent can unilaterally initiate complete replication regardless of its governance programming status, because complete agent assembly requires multi-agent coordination that is itself subject to resource-budget oversight. The partial replication architecture is assessed as having extrapolated viability for deep-space application while remaining theoretically grounded in its structural governance properties.

4.4 The Checkpoint Horizon: Theoretical Development

The Checkpoint Horizon is the paper's principal original theoretical contribution, and its analytical development requires distinguishing it carefully from the existing relay-node literature on which it builds. Current deep-space relay network concepts treat relay nodes exclusively as communication infrastructure — signal amplification and retransmission waypoints whose function is to extend the communication range of the primary vessel (Fraire et al., 2018). The Checkpoint Horizon concept reconceptualizes each relay node's informational function along two additional dimensions.

The first dimension is archival: at the moment of deployment, each relay node receives and stores a complete snapshot of the mission's state — navigational position and trajectory, systems health telemetry, scientific data accumulated to that point, AI operational parameters, swarm population status, and resource inventory. This information is written to redundant radiation-hardened storage aboard the relay node and is preserved independently of the vessel's subsequent operational status. If the vessel suffers a catastrophic failure after deploying a relay node, that node's archived snapshot represents the last fully recoverable statement of the mission's progress — the furthest point from which its trajectory and accumulated data can be reconstructed by Earth. The relay node is, in this sense, a recoverable save point: a persistent record of mission existence and status that survives independently of the mission's own survival.

The second dimension is coordinative: each relay node's archived snapshot, combined with the known physics of the vessel's propulsion and trajectory, enables Earth — at any future time, including centuries after the mission's departure — to reconstruct the vessel's probable position, operational status, and accumulated scientific data with a precision that improves as Earth's own computational and observational capabilities advance. A future Earth-launched mission, departing decades or centuries after the original Ariadne vessel, can use the relay network's sequence of archived snapshots as a positional and informational guide — navigating not into an unknown void but toward a documented trajectory with known waypoints, each represented by a relay node whose stored state provides a reconstruction basis for the preceding mission segment. The Thread, in this formulation, transforms the interstellar void from a space of total informational opacity into a navigated corridor with documented history.

The Checkpoint Horizon concept is assessed as theoretically grounded in its first dimension — the archival function — because it rests on established principles of DTN store-and-forward architecture (Fraire et al., 2018) and redundant radiation-hardened storage (Laizans et al., 2014), extrapolated to relay nodes of greater storage capacity and operational duration than currently deployed. Its second dimension — the coordinative function enabling multi-generational mission sequencing — is assessed as speculative in its full realization, because it depends on future Earth-launched missions being dispatched, on the relay nodes remaining operational across the required timescales, and on Earth's retention of the institutional knowledge and technical capability to interpret and act on the archived data.

4.5 Challenge Assessment: Principal Unresolved Problems

Three principal challenges identified in the literature require direct analytical engagement. The first is runaway replication risk (Ellery, 2022). The framework's response — layered governance combining partial replication architecture, apoptosis-analogue decommissioning protocols, population-density feedback, and hard resource-budget limits — addresses this risk through redundant independent mechanisms rather than relying on any single constraint. The theoretical robustness of this layered approach is supported by the principle, established across the fault tolerance and biosafety engineering literatures, that independent redundant safeguards produce failure probabilities that are multiplicatively rather than additively lower than any single safeguard (Malik, 2025). The framework does not claim to eliminate replication risk; it claims to reduce it to a level commensurate with other long-duration mission risks through a governance architecture whose components are individually theoretically grounded.

The second challenge is interstellar relay communication bit-rate constraints. The literature establishes that laser communication relay networks can theoretically sustain meaningful data rates across light-year distances, but that relay node power budgets impose severe constraints on retransmission capacity that reduce achievable bit rates far below what current mission communication architectures assume (Hippke, 2017, 2020). The framework's response is to reframe the communication requirement: the Ariadne architecture does not require real-time high-bandwidth communication. The relay network's primary functions — delivering periodic machine learning updates to the vessel, transmitting mission state snapshots from the vessel to relay nodes, and preserving archived snapshots for Earth's future access — are each compatible with low-bandwidth, high-latency, store-and-forward transmission. The reduction in bandwidth requirement substantially relaxes the relay node power budget constraint and renders the system viable within power budgets achievable by foreseeable radioisotope or compact fission power systems.

The third challenge is hardware resupply. Unlike crewed missions, which can carry manufacturing feedstocks and replacement hardware for the duration of the transit, an interstellar vessel operating for centuries will eventually exhaust materials that cannot be replenished until a target system with accessible resources is reached. The framework's response draws on the ISRU literature: the vessel is designed to harvest interstellar medium dust and gas — whose density, while extremely low, is non-zero across all interstellar trajectories (Long, 2023; Loup, 2017) — through electromagnetic collection systems, supplementing onboard feedstock reserves across the transit. The theoretical viability of interstellar medium resource harvesting is established in principle (Lacki, 2021; Karimov et al., 2022) but not validated for the material compositions and collection efficiencies required by the framework's manufacturing demands, placing this response in the extrapolated epistemic register. The framework treats feedstock exhaustion as a mission-duration constraint that defines a practical operational horizon rather than a design-time catastrophic failure.

5. Discussion of Findings

5.1 Reframing the Definition of Exploration

The analytical synthesis of Section 4 has implications that extend beyond the engineering assessment of any individual pillar. Taken together, the three-pillar architecture and the Checkpoint Horizon concept constitute an implicit argument about what exploration means when the explorer is not human; this argument deserves explicit treatment because it bears directly on how the framework should be evaluated, funded, and developed as a research program.

The conventional definition of space exploration is implicitly anthropocentric: exploration is the extension of human presence, human perception, and human experience into new territory (Lester & Thronson, 2011; Wingerden & Vigneswaran, 2024). On this definition, a mission that reaches another star system but carries no human being has not explored that system in any meaningful sense — it has merely sent an instrument. The Ariadne framework challenges this definition at its foundation. If the vessel that reaches the target system is itself an intelligence — capable of perception, decision-making, adaptive response, and the accumulation and transmission of knowledge — then the question of whether a biological human is present aboard is a question about the substrate of the intelligence rather than about the presence of intelligence itself. The exploration is occurring; the explorer happens to be distributed across a network of silicon and metal rather than concentrated in a body of carbon and water.

This reframing is not merely philosophical. It has direct implications for how mission success criteria are defined, how scientific return is evaluated relative to mission cost, and how public and institutional support for interstellar exploration programs is framed. A program oriented toward crewed interstellar exploration faces an almost insuperable political and ethical challenge in justifying the deliberate exposure of human beings — or their descendants — to the extreme risks of a centuries-long transit (Neukart, 2024; Rajput et al., 2023). A program oriented toward intelligence-extended exploration, in which the mission's continuity and adaptability are guaranteed by engineering rather than by biological resilience, faces a substantially different and more tractable public justification challenge.

5.2 Relationship to the Fermi Paradox

The Fermi Paradox dimension of the self-replicating probe literature, reviewed in Section 2.1, acquires a specific interpretive significance when read against the Ariadne framework. The standard Fermi Paradox formulation asks why, if self-replicating probes are feasible and any technological civilization has deployed them, we do not observe their presence in the solar system or detect their signals (Forgan, 2019; Kowald, 2016). The Ariadne framework's governed, communication-anchored variant of probe deployment offers a partial response to this question that has not been explicitly developed in the literature.

An Ariadne-type probe network — in which the vessel deploys relay nodes along its trajectory but does not replicate itself into new independent vessels directed toward additional targets — would produce a galactic signature fundamentally different from the exponentially diffusing von Neumann probe network that Tipler (1980) and subsequent authors have used in Fermi Paradox reasoning. The Ariadne Thread extends linearly along a single trajectory; it does not branch, replicate toward new targets, or fill the galaxy. Its signal would be a thin, directed line of relay nodes rather than the isotropic galactic presence that Fermi Paradox arguments typically require. A civilization deploying Ariadne-type missions would leave a galactic footprint that is detectable only along the precise trajectory of each individual mission — an observational needle in a galactic haystack that current SETI instrumentation would not reliably detect even if multiple such missions had been launched by other civilizations.

This observation does not resolve the Fermi Paradox, and the paper does not claim that it does. It does, however, establish that the Ariadne framework's design choices — in particular the deliberate constraint of replication to within-vessel renewal rather than between-vessel proliferation — have implications for the observable signatures of technological civilization activity that the existing literature has not examined in this specific form.

5.3 Ethical and Governance Dimensions

The ethical and governance dimensions of Project Ariadne operate at two distinct levels that require separate treatment. The first level concerns the deployment decision itself: the governance frameworks applicable to launching a centuries-long autonomous mission whose operational decisions no living human being can supervise and whose physical presence will eventually be beyond any jurisdiction currently recognized in international space law. The Outer Space Treaty of 1967 and its successor instruments were developed in a context of crewed and near-Earth robotic missions; they provide no adequate framework for an autonomous, self-replicating vessel operating beyond the solar system across timescales that exceed any institutional continuity that current nation-states or international bodies can plausibly guarantee (Bratu et al., 2021; Yap & Truffer, 2022).

This governance gap is not a reason to abandon the framework, but it is a reason to recognize that the development of Project Ariadne as an operational program would require the parallel development of new international governance instruments addressing, at a minimum, three specific questions: the legal status of an autonomous self-replicating vessel beyond national jurisdiction; the liability and responsibility framework for a mission whose operational decisions are made by an artificial intelligence without human oversight; and the conditions under which a future civilization — potentially including one at the target stellar system, if one exists — would have legitimate grounds to disable or redirect the mission.

The second level concerns the informational and cultural authority embedded in the mission. An Ariadne vessel carries, at minimum, the accumulated scientific and cultural knowledge encoded in its AI architecture's training data and operational parameters — a representation of human civilization at a specific historical moment. The question of whose knowledge, whose values, and whose priorities are encoded in the vessel's intelligence — and who has the authority to make those encoding decisions — is an ethical question of the first order that the framework's technical architecture does not resolve. It is analogous to, but more permanent and consequential than, the questions raised by the decisions on cultural content embodied in the Voyager Golden Record (Capova, 2021), and it requires a deliberative process of commensurate seriousness.

5.4 Implications for Near-Term Research and Mission Architecture

The Ariadne framework's analytical assessment identifies a specific set of near-term research priorities whose advancement would most directly reduce the framework's principal uncertainties. The governed replication pillar's most pressing research need is the experimental validation of apoptosis-analogue governance in self-replicating robotic systems at scales and durations sufficient to test governance fidelity across multiple replication generations — work that is tractable within current robotics and manufacturing research capabilities and does not require space deployment (Wyder et al., 2024). The AI-as-ship pillar's most pressing research need is the development of architectural standards for autonomous AI systems designed for decadal-scale unsupervised operation, including formal specifications for modular update compatibility and degradation-tolerant decision-making. The relay network pillar's most pressing research need is the engineering of relay node power systems and storage media with operational lifetimes measured in decades rather than years.

Collectively, these research needs define a program that is ambitious but not discontinuous with current scientific and engineering capability (Shahzad et al., 2023; Volpe et al., 2024). The Ariadne framework does not require a technological revolution in any single domain; it requires sustained incremental progress in four domains — autonomous AI, swarm robotics, self-replication governance, and deep-space communication engineering — whose convergence, rather than any single breakthrough, produces the integrated capability the framework requires. This convergence profile is, if anything, more favorable from a research planning perspective than a framework dependent on a single speculative technological leap, because it distributes risk across multiple independent research tracks and allows partial progress in each track to generate mission-relevant capability even before the full framework is realizable.

6. Conclusions and Recommendations

6.1 Summary of the Framework and Its Theoretical Contributions

This paper has proposed, developed, and analytically evaluated Project Ariadne — a three-pillar conceptual framework for interstellar exploration in which mission intelligence, physical continuity, and informational persistence are distributed across machine systems rather than being dependent on human biological survival. The framework's first pillar, the AI-as-ship architecture, reconceptualizes the spacecraft itself as a distributed intelligence rather than a vessel carrying intelligence as a separate system, drawing on established principles of distributed AI (Goonatilleke & Hettige, 2022; Wei & Liu, 2024) and autonomous spacecraft design (Nesnas et al., 2021) to propose an architecture that is fault-tolerant, adaptable, and remotely upgradeable through the relay network. The second pillar, governed swarm replication, addresses the physical continuity requirement of centuries-long operation through a population-controlled model of robotic self-renewal governed by layered mechanisms — partial replication architecture (Abdel-Rahman et al., 2022), apoptosis-analogue decommissioning (Sterritt, 2011), population-density feedback (Bradshaw & Herrando-Perez, 2023; Mayya et al., 2019), and resource-budget constraints (Figueiredo et al., 2025; Wu & Durfee, 2010) — each individually grounded in existing theoretical literature and collectively providing redundant protection against the runaway replication failure mode. The third pillar, the Ariadne Thread, establishes a progressive relay-node network deployed along the vessel's trajectory that sustains bi-directional communication between the mission and Earth across any distance, functioning as both an update delivery channel and a persistent informational spine connecting the mission to its origin civilization.

The paper's principal original theoretical contribution is the Checkpoint Horizon concept, which reconceptualizes relay nodes from passive signal repeaters to active mission state archives — recoverable save points that preserve a complete snapshot of mission status at each deployment interval, independently of the vessel's subsequent operational fate, and that enable future Earth-launched missions to reconstruct, locate, and coordinate with the original mission regardless of elapsed time. The Checkpoint Horizon extends current DTN and relay-node literature into a dimension — informational archival and asynchronous multi-generational coordination — that the existing literature has not previously formalized (Burleigh et al., 2003; Elewaily et al., 2023; Fraire et al., 2018).

6.2 Theoretical Defensibility and Honest Limitations

The framework's three pillars are each theoretically grounded in established scientific and engineering literature (Hein, 2016; Hein & Baxter, 2018; Wang, 2025), extrapolated to a scale and duration that the literature has not yet directly addressed (Hein & Baxter, 2018). The paper has distinguished consistently among established, theoretically grounded, extrapolated, and speculative epistemic registers in its evaluative claims, and those distinctions stand as the paper's candid accounting of what the framework currently is: a theoretically coherent and literature-supported conceptual architecture whose operational viability requires sustained research across four disciplines before it approaches engineering specification. Project Ariadne is not a mission proposal; it is a framework that makes a mission of this type conceivable and that identifies the research agenda required to make it achievable.

6.3 Recommendations

Four recommendations follow from the paper's analysis, directed at distinct communities of relevant researchers and institutions.

First, the autonomous systems and AI safety research communities are encouraged to develop architectural standards and formal design specifications for AI systems intended for unsupervised decadal-scale operation — a design category that currently lacks both formal standards and validated reference architectures, and whose development is a prerequisite not only for the Ariadne framework but for any deep-space mission architecture of comparable operational independence.

Second, the swarm robotics and advanced manufacturing research communities are encouraged to prioritize experimental programs examining governance fidelity across multiple replication generations in self-replicating robotic systems. The theoretical mechanisms for governed replication are established; their experimental validation in hardware systems operating across realistic timescales and under realistic environmental stress is the critical missing empirical foundation for the framework's second pillar.

Third, the deep-space communication engineering community is encouraged to examine relay node design under the dual-function requirements of the Checkpoint Horizon concept — that is, relay nodes engineered simultaneously for signal retransmission and for long-duration mission state archival — as a specific design target distinct from relay nodes optimized exclusively for communication throughput. This dual-function design problem has not been formally addressed in the literature and represents a tractable near-term engineering research question.

Fourth, the international space law and governance community is encouraged to initiate foundational work on the legal and ethical frameworks applicable to autonomous, self-replicating missions operating beyond Solar System jurisdiction across multi-century timescales. The technical development of such missions will outpace governance framework development if the latter does not begin substantially in advance of the former (Pohler et al., 2024).

6.4 Closing Argument

The history of exploration is, at its deepest level, the history of humanity's successive expansions of the definition of the reachable. Project Ariadne proposes that the next expansion requires not a larger ship or a longer journey, but a reconception of who — or what — does the reaching. Intelligence, extended outward through a self-renewing vessel, anchored to its origin by a thread of relay nodes, and capable of receiving the accumulated wisdom of an advancing civilization across the centuries of its transit, is not a diminished substitute for human exploration. It is exploration made permanent, recoverable, and indefinitely extensible — one dropped anchor at a time.

References

Abdel-Rahman, A., Cameron, C., Jenett, B., Smith, M., & Gershenfeld, N. (2022). Self-replicating hierarchical modular robotic swarms. Communications Engineering, 1(1). https://doi.org/10.1038/s44172-022-00034-3

Alhilal, A., Braud, T., & Hui, P. (2019). The sky is NOT the limit anymore: Future architecture of the interplanetary Internet. IEEE Aerospace and Electronic Systems Magazine, 34(8), 22. https://doi.org/10.1109/maes.2019.2927897

Andersson, J., Hojcková, K., & Sandén, B. A. (2023). On the functional and structural scope of technological innovation systems — A literature review with conceptual suggestions. Environmental Innovation and Societal Transitions, 49, 100786. https://doi.org/10.1016/j.eist.2023.100786

Arafeh, M., Otrok, H., Ould-Slimane, H., Mourad, A., Talhi, C., & Damiani, E. (2023). ModularFed: Leveraging modularity in federated learning frameworks. Internet of Things, 22, 100694. https://doi.org/10.1016/j.iot.2023.100694

Araniti, G., Bezirgiannidis, N., Birrane, E., Bisio, I., Burleigh, S., Caini, C., Feldmann, M., Marchese, M., Segui, J., & Suzuki, K. (2015). Contact graph routing in DTN space networks: Overview, enhancements and performance. IEEE Communications Magazine, 53(3), 38. https://doi.org/10.1109/mcom.2015.7060480

Armstrong, S., & Sandberg, A. (2013). Eternity in six hours: Intergalactic spreading of intelligent life and sharpening the Fermi paradox. Acta Astronautica, 89, 1–13.

Azami, M., Kazemi, Z., Moazen, S., Dubé, M., Potvin, M.-J., & Skonieczny, K. (2024). A comprehensive review of lunar-based manufacturing and construction. arXiv. https://doi.org/10.48550/arxiv.2408.05823

Balachandran, A., Harasha, N., & Lynch, N. (2024). Swarm algorithms for dynamic task allocation in unknown environments. arXiv. https://doi.org/10.48550/arxiv.2409.09550

Baratoux, D., Aoudjehane, H. C., Gibson, R. L., et al. (2017). The state of planetary and space sciences in Africa. Eos. https://doi.org/10.1029/2017eo075833

Barlow, M. T. (2013). Galactic exploration by directed self-replicating probes, and its implications for the Fermi paradox. International Journal of Astrobiology, 12(1), 63–68.

Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). Swarm intelligence: From natural to artificial systems. Oxford University Press.

Borgue, O., & Hein, A. M. (2021). Near-term self-replicating probes: A concept design. Acta Astronautica, 187, 546. https://doi.org/10.1016/j.actaastro.2021.03.004

Bossens, D. M., & Tarapore, D. (2020). QED: Using quality-environment-diversity to evolve resilient robot swarms. IEEE Transactions on Evolutionary Computation, 25(2), 346. https://doi.org/10.1109/tevc.2020.3036578

Bradshaw, C. J. A., & Herrando-Pérez, S. (2023). Logistic-growth models measuring density feedback are sensitive to population declines, but not fluctuating carrying capacity. Ecology and Evolution, 13(4). https://doi.org/10.1002/ece3.10010

Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. Swarm Intelligence, 7(1), 1–41.

Bratu, I., Lodder, A. R., & Linden, M. van der. (2021). Autonomous space objects and international space law: Navigating the liability gap. Indonesian Journal of International Law, 18(3). https://doi.org/10.17304/ijil.vol18.3.818

Buchanan, E., Alden, K., Pomfret, A., Timmis, J., & Tyrrell, A. M. (2023). A study of error diversity in robotic swarms for task partitioning in foraging tasks. Frontiers in Robotics and AI, 9. https://doi.org/10.3389/frobt.2022.904341

Burleigh, S., Hooke, A. J., Torgerson, L., Fall, K., Cerf, V. G., Durst, B., Scott, K., & Weiss, H. (2003). Delay-tolerant networking: An approach to interplanetary Internet. IEEE Communications Magazine, 41(6), 128. https://doi.org/10.1109/mcom.2003.1204759

Caini, C., Cruickshank, H., Farrell, S., Marchese, M., Pentikousis, K., Poti, R., Rinaldi, R., & Tsaoussidis, V. (2011). Delay- and disruption-tolerant networking (DTN): An alternative solution for future satellite networking applications. Proceedings of the IEEE, 99(11), 1980–1997. https://doi.org/10.1109/jproc.2011.2158378

Capova, K. A. (2021). Introducing humans to the extraterrestrials: The pioneering missions of the Pioneer and Voyager probes. Frontiers in Human Dynamics, 3. https://doi.org/10.3389/fhumd.2021.714616

Cardion, A. L. (2015). Ideal biological characteristics for long-duration manned space travel. Journal of the British Interplanetary Society, 68, 31.

Cerf, V., Burleigh, S., Hooke, A., Torgerson, L., Durst, R., Scott, K., Fall, K., & Weiss, H. (2007). Delay-tolerant networking architecture (RFC 4838). Internet Engineering Task Force.

Chang, H., Myeong, J., Kim, S., & Mazzoleni, A. P. (2018). Marsbee: Swarm of flapping wing flyers for enhanced Mars exploration. NASA NIAC Phase I Final Report.

Chen, H., Ornik, M., & Ho, K. (2021). Space exploration architecture and design framework for commercialization. Journal of Spacecraft and Rockets, 59(2), 538. https://doi.org/10.2514/1.a35077

Chen, X., Zhang, H., & Liu, Y. (2024). Swarm robotics for in-situ resource utilization in lunar construction: A review. MDPI Applied Sciences, 14(3), 1–22.

Chien, S., Sherwood, R., Tran, D., Cichy, B., Rabideau, G., & Castano, R. (2004). Onboard autonomy on the Earth Observing One mission. https://doi.org/10.2514/6.2004-6515

Childress, S. D., Williams, T. C., & Francisco, D. (2023). NASA space flight human-system standard: Enabling human spaceflight missions by supporting astronaut health, safety, and performance. npj Microgravity, 9(1). https://doi.org/10.1038/s41526-023-00275-2

Communication Papers of the 17th Conference on Computer Science and Intelligence Systems. (2022). Annals of Computer Science and Information Systems, 32. https://doi.org/10.15439/978-83-965897-4-3

Curtis, S. A., Mica, J., Nuth, J., Marr, G., Wessner, M., & Vranish, J. (2003). ANTS (Autonomous Nano-Technology Swarm): An artificial life approach to asteroid belt exploration. In Proceedings of the AIAA Space 2003 Conference.

Cussat-Blanc, S., Harrington, K., & Banzhaf, W. (2019). Artificial gene regulatory networks — A review. Artificial Life, 24(4), 296. https://doi.org/10.1162/artl_a_00267

Devi, K. V. R., Smitha, B. S., Lakhanpal, S., Kalra, R., Sethi, V. A., & Thajil, S. K. (2024). A review: Swarm robotics: Cooperative control in multi-agent systems. E3S Web of Conferences, 505, 3013. https://doi.org/10.1051/e3sconf/202450503013

Dias, C. S., Trivedi, M., Volpe, G., Araújo, N. A. M., & Volpe, G. (2023). Environmental memory boosts group formation of clueless individuals. Nature Communications, 14(1). https://doi.org/10.1038/s41467-023-43099-0

Dorigo, M., & Birattari, M. (2007). Swarm intelligence. Scholarpedia, 2(9), 1462.

Dorigo, M., Floreano, D., Gambardella, L. M., Mondada, F., Nolfi, S., Baaboura, T., Birattari, M., Bonani, M., Brambilla, M., Brutschy, A., & Burnier, D. (2014). Swarmanoid: A novel concept for the study of heterogeneous robotic swarms. IEEE Robotics & Automation Magazine, 20(4), 60–71.

Drexler, K. E. (1986). Engines of creation: The coming era of nanotechnology. Anchor Press.

Duarte, M., Gomes, J., Costa, V., Rodrigues, T., Silva, J. F., Lobo, V., Marques, M. M., Oliveira, S., & Christensen, A. L. (2016). Application of swarm robotics systems to marine environmental monitoring. OCEANS 2016 - Shanghai, 1. https://doi.org/10.1109/oceansap.2016.7485429

Elewaily, D. I., Ali, H., Saleh, A. I., & Abdelsalam, M. M. (2023). Delay/disruption-tolerant networking-based the integrated deep-space relay network: State-of-the-art. Ad Hoc Networks, 152, 103307. https://doi.org/10.1016/j.adhoc.2023.103307

Ellery, A. (2022a). Curbing the fruitfulness of self-replicating machines. International Journal of Astrobiology, 21(4), 243. https://doi.org/10.1017/s1473550422000246

Ellery, A. (2022b). Self-replicating probes are imminent — Implications for SETI. International Journal of Astrobiology, 21(4), 212. https://doi.org/10.1017/s1473550422000234

Ellery, A. (2025). The state of hybrid artificial intelligence for interstellar missions. Progress in Aerospace Sciences, 156, 101100. https://doi.org/10.1016/j.paerosci.2025.101100

Elmore, S. (2007). Apoptosis: A review of programmed cell death. Toxicologic Pathology, 35(4), 495–516.

Estlin, T., Bornstein, B., Gaines, D., Anderson, R. C., Thompson, D. R., Burl, M. C., Castano, R., & Judd, M. (2012). AEGIS automated targeting for MER opportunity rover.

Fayyaz, M., & Vladimirova, T. (2016). Survey and future directions of fault-tolerant distributed computing on board spacecraft. Advances in Space Research, 58(11), 2352. https://doi.org/10.1016/j.asr.2016.08.017

Figueiredo, R. P. de, Limberg, C., Jamone, L., & Bernardino, A. (2025). Editorial: Computer vision mechanisms for resource-constrained robotics applications. Frontiers in Robotics and AI, 12, 1680098. https://doi.org/10.3389/frobt.2025.1680098

Forgan, D. H. (2019). Predator-prey behaviour in self-replicating interstellar probes. International Journal of Astrobiology, 18(6), 552. https://doi.org/10.1017/s1473550419000053

Fraire, J. A., Feldmann, M., Walter, F., Fantino, E., & Burleigh, S. (2018). Networking in interstellar dimensions: Communicating with TRAPPIST-1. IEEE Transactions on Aerospace and Electronic Systems, 55(4), 1656. https://doi.org/10.1109/taes.2018.2874149

Freitas, R. A., Jr. (1980). A self-reproducing interstellar probe. Journal of the British Interplanetary Society, 33, 251–264.

Goonatilleke, S. T., & Hettige, B. (2022). Past, present and future trends in multi-agent system technology. Journal Européen des Systèmes Automatisés, 55(6), 723. https://doi.org/10.18280/jesa.550604

Harkins, A. M. (1978). Futurism in man: Humanism, social technology, and survival. In De Gruyter eBooks (p. 125). De Gruyter. https://doi.org/10.1515/9783110800067.125

Hein, A. M. (2016). Artificial intelligence probes for interstellar exploration and colonization. arXiv. https://doi.org/10.48550/arxiv.1612.08733

Hein, A. M., & Baxter, S. (2018). Artificial intelligence for interstellar travel. arXiv. https://doi.org/10.48550/arxiv.1811.06526

Hein, A. M., & Baxter, S. (2023). Artificial intelligence for interstellar travel. In K. F. Long (Ed.), Proceedings of the Initiative for Interstellar Studies Symposium.

Hein, A. M., Smith, C. M., Marin, F., & Staats, K. (2020). World ships: Feasibility and rationale. HAL. https://hal.science/hal-02561266

Hinchey, M., Sterritt, R., & Rouff, C. (2007). Swarms and swarm intelligence. Computer, 40(4), 111. https://doi.org/10.1109/mc.2007.144

Hippke, M. (2017). Interstellar communication I: Maximized data rate for lightweight space-probes. arXiv. https://doi.org/10.48550/arxiv.1706.03795

Hippke, M. (2020). Interstellar communication network II: Deep space nodes with gravitational lensing. arXiv. https://doi.org/10.48550/arxiv.2009.01866

Hippke, M. (2021). Interstellar communication network I: Overview and assumptions. The Astronomical Journal, 162(1), 1–10.

Jones, H., & Kliss, M. (2009). Exploration life support technology challenges for the crew exploration vehicle and future human missions. Advances in Space Research, 45(7), 917. https://doi.org/10.1016/j.asr.2009.10.018

Karimov, A. R., Yakovlev, O. V., & Murad, P. (2022). Using interplanetary medium as propellant for plasma thruster propulsion. arXiv. https://doi.org/10.48550/arxiv.1808.03334

Klee, R. (2018). The importance of being sufficiently realistic: A reply to Milan Cirkovic. International Journal of Astrobiology, 18(1), 10. https://doi.org/10.1017/s1473550417000404

Konecny, J., McMahan, H. B., Ramage, D., & Richtarik, P. (2016). Federated optimization: Distributed machine learning for mobile devices. arXiv preprint arXiv:1610.02527.

Kowald, A. (2016). Why is there no von Neumann probe on Ceres? Error catastrophe can explain the Fermi-Hart paradox. arXiv. https://doi.org/10.48550/arxiv.1605.02169

Lacki, B. C. (2021). Galactic traversability: A new concept for extragalactic SETI. International Journal of Astrobiology, 20(5), 359. https://doi.org/10.1017/s1473550421000252

Laizans, K., Sünter, I., Zalite, K., Kuuste, H., Valgur, M., Tarbe, K., Allik, V., Olentsenko, G., Laes, P., Lätt, S., & Noorma, M. (2014). Design of the fault tolerant command and data handling subsystem for ESTCube-1. Proceedings of the Estonian Academy of Sciences, 63, 222. https://doi.org/10.3176/proc.2014.2s.03

Lange, K., Fontana, F., Rossi, F., Varile, M., & Apruzzese, G. (2024). Machine learning in space: Surveying the robustness of on-board ML models to radiation. arXiv. https://doi.org/10.48550/arxiv.2405.02642

Langeveld, M., Sheratt, T. N., & Skyrms, B. (2022). Predator-prey dynamics and mutation in self-replicating probe populations. PLoS Computational Biology, 18(4), e1010062.

Langton, C. G. (1984). Self-reproduction in cellular automata. Physica D: Nonlinear Phenomena, 10, 135. https://doi.org/10.1016/0167-2789(84)90256-2

Lesser, V. R. (1999). Cooperative multiagent systems: A personal view of the state of the art. IEEE Transactions on Knowledge and Data Engineering, 11(1), 133–142.

Lester, D. F., & Thronson, H. A. (2011). Human space exploration and human spaceflight: Latency and the cognitive scale of the universe. Space Policy, 27(2), 89. https://doi.org/10.1016/j.spacepol.2011.02.002

Lo, S. K., Lu, Q., Wang, C., Paik, H.-Y., & Zhu, L. (2020). A systematic literature review on federated machine learning: From a software engineering perspective. arXiv. https://arxiv.org/abs/2007.11354

Long, K. F. (2023). Calculations of particle bombardment due to dust and charged particles in the ISM on the Project Starshot gram-scale interstellar probe. Journal of the British Interplanetary Society, 76(8), 262. https://doi.org/10.59332/jbis-076-08-0262

Loup, F. (2017). The analysis of Chris Van Den Broeck applied to the Natario warp drive spacetime using the original Alcubierre shape function. HAL. https://hal.science/hal-01456561

Low, K. H., Leow, W. K., & Ang, M. H. (2006). Autonomic mobile sensor network with self-coordinated task allocation and execution. IEEE Transactions on Systems, Man, and Cybernetics — Part C, 36(3), 315. https://doi.org/10.1109/tsmcc.2006.871590

Maccone, C. (2009). Deep space flight and communications: Exploiting the Sun as a gravitational lens. Springer.

Malik, A. (2025). Evaluation of isolation design flow (IDF) for single chip cryptography (SCC) application. arXiv. https://doi.org/10.48550/arxiv.2501.16217

Marcy, G. W., Isaacson, H., Siemion, A., & Lebofsky, M. (2022). Engineering an interstellar communications network. International Journal of Astrobiology, 21(5), 309–323.

Marin, F., & Beluffi, C. (2018). Computing the minimum crew for a multigenerational trip toward Proxima Centauri b. Journal of the British Interplanetary Society, 71(2), 45–52.

Marin, F., Beluffi, C., & Fischer, F. J. (2022). Genetic evolution of a multi-generational population in the context of interstellar space travels — Part I: Genetic evolution under the neutral selection hypothesis. arXiv. https://doi.org/10.48550/arxiv.2102.01508

Matloff, G. L. (2022). Von Neumann probes: Rationale, propulsion, interstellar transfer timing. Acta Astronautica, 193, 609–613.

Mayya, S., Pierpaoli, P., & Egerstedt, M. (2019). Voluntary retreat for decentralized interference reduction in robot swarms. Proceedings of the 2019 International Conference on Robotics and Automation (ICRA), 9667. https://doi.org/10.1109/icra.2019.8794124

McMullin, B. (2000). John von Neumann and the evolutionary growth of complexity: Looking backwards, looking forwards. In The MIT Press eBooks (p. 467). https://doi.org/10.7551/mitpress/1432.003.0065

Naylor, B., Read, M., Timmis, J., & Tyrell, A. (2014). The relay chain: A scalable dynamic communication link between an exploratory underwater shoal and a surface vehicle. https://doi.org/10.7551/978-0-262-32621-6-ch047

Nesnas, I., Fesq, L., & Volpe, R. (2021). Autonomy for space robots: Past, present, and future. Current Robotics Reports, 2(3), 251. https://doi.org/10.1007/s43154-021-00057-2

Neukart, F. (2024). Toward the stars: Technological, ethical, and sociopolitical dimensions of interstellar exploration. arXiv. https://doi.org/10.48550/arxiv.2402.15536

Perryman, M. A. C. (2012). The history of astrometry. The European Physical Journal H, 37(5), 745. https://doi.org/10.1140/epjh/e2012-30039-4

Pesavento, U. (1995). An implementation of von Neumann's self-reproducing machine. Artificial Life, 2(4), 337. https://doi.org/10.1162/artl.1995.2.337

Pohler, L. D., Diepold, K., & Wallach, W. (2024). A practical multilevel governance framework for autonomous and intelligent systems. https://doi.org/10.48550/ARXIV.2404.13719

Rafi, M. A., & Alavi, A. (2017). Debate on human aging and lifespan. Bioimpacts, 7(3), 135. https://doi.org/10.15171/bi.2017.16

Rajput, S., Mayor, I., Diamond, M., Rosenberg, M., Cole, V., Bhakare, N. S., Aziz, O., & Wilder, A. (2023). Medical ethics of long-duration spaceflight. npj Microgravity, 9(1). https://doi.org/10.1038/s41526-023-00333-9

Roy, S., Biswas, S., & Chaudhuri, S. S. (2014). Nature-inspired swarm intelligence and its applications. International Journal of Modern Education and Computer Science, 6(12), 55. https://doi.org/10.5815/ijmecs.2014.12.08

Sadownik, J. W., Kosikova, T., & Philp, D. (2017). Generating system-level responses from a network of simple synthetic replicators. Journal of the American Chemical Society, 139(48), 17565. https://doi.org/10.1021/jacs.7b09735

Sagan, C., & Newman, W. I. (1983). The solipsist approach to extraterrestrial intelligence. Quarterly Journal of the Royal Astronomical Society, 24, 113–121.

Scharf, D. P., Hadaegh, F. Y., & Ploen, S. R. (2004). A survey of spacecraft formation flying guidance and control. In Proceedings of the American Control Conference (Vol. 2, pp. 1733–1739).

Schwartz, J. T., Neumann, J. von, & Burks, A. W. (1967). Theory of self-reproducing automata. Mathematics of Computation, 21(100), 745. https://doi.org/10.2307/2005041

Scott, K., & Burleigh, S. (2007). Bundle protocol specification (RFC 5050). Internet Engineering Task Force.

Sebastián, A., Acedo, L., & Moraño, J.-A. (2022). An orbital model for the Parker Solar Probe mission: Classical vs relativistic effects. Advances in Space Research, 70(3), 842. https://doi.org/10.1016/j.asr.2022.05.037

Shahzad, M. M., Saeed, Z., Akhtar, A., Munawar, H., Yousaf, M. H., Baloach, N. K., & Hussain, F. (2023). A review of swarm robotics in a nutshell. Drones, 7(4), 269. https://doi.org/10.3390/drones7040269

Simões, P. C., Moreira, A. C., & Dias, C. M. M. (2022). The 'endless perspective' to university-industry-government relations. Triple Helix Journal, 9(3), 247. https://doi.org/10.1163/21971927-bja10034

Smith, C. M. (2014). Estimation of a genetically viable population for multigenerational interstellar voyaging: Review and data for project Hyperion. Acta Astronautica, 97, 16–29.

Steinberg, M., Stack, J. R., & Paluszkiewicz, T. (2016). Long duration autonomy for maritime systems: Challenges and opportunities. Autonomous Robots, 40(7), 1119. https://doi.org/10.1007/s10514-016-9582-0

Sterritt, R. (2011). Apoptotic robotics: Programmed death by default. Proceedings of EASE 2011, 82, 107. https://doi.org/10.1109/ease.2011.21

Tan, J., Melkoumian, N., Harvey, D., & Akmeliawati, R. (2024). Lunarminer framework for nature-inspired swarm robotics in lunar water ice extraction. Biomimetics, 9(11), 680. https://doi.org/10.3390/biomimetics9110680

Timmons, T., Bailet, G., Beeley, J., & McInnes, C. R. (2021). Mars atmospheric characterization with a ChipSat swarm. Journal of Spacecraft and Rockets, 58(5), 1453. https://doi.org/10.2514/1.a34970

Tipler, F. J. (1980). Extraterrestrial intelligent beings do not exist. Quarterly Journal of the Royal Astronomical Society, 21, 267–281.

Trindade, M. G. (2021). Assessment of edge machine-learning systems under radiation-induced effects [Doctoral thesis, HAL]. https://theses.hal.science/tel-03525899

Truszkowski, W., Hinchey, M., Rash, J. L., & Rouff, C. (2006). Autonomous and autonomic systems: A paradigm for future space exploration missions. IEEE Transactions on Systems, Man, and Cybernetics — Part C, 36(3), 279. https://doi.org/10.1109/tsmcc.2006.871600

Varadharajan, V. S., St-Onge, D., Adams, B., & Beltrame, G. (2020). Swarm relays: Distributed self-healing ground-and-air connectivity chains. IEEE Robotics and Automation Letters, 5(4), 5347. https://doi.org/10.1109/lra.2020.3006793

Varpio, L., Paradis, E., Uijtdehaage, S., & Young, M. (2020). The distinctions between theory, theoretical framework, and conceptual framework. Academic Medicine, 95(7), 989–994.

Verma, V., Maimone, M., Gaines, D., Francis, R., Estlin, T., Kuhn, S., Rabideau, G., Chien, S., McHenry, M., Graser, E., Rankin, A., & Thiel, E. (2023). Autonomous robotics is driving Perseverance rover's progress on Mars. Science Robotics, 8(80). https://doi.org/10.1126/scirobotics.adi3099

Volpe, G., Araújo, N. A. M., Guix, M., et al. (2024). Roadmap for animate matter. arXiv. https://doi.org/10.48550/arxiv.2407.10623

von Neumann, J. (1966). Theory of self-reproducing automata (A. W. Burks, Ed.). University of Illinois Press.

Wang, Z. (2025). Space AI: Leveraging artificial intelligence for space to improve life on Earth. arXiv. https://doi.org/10.48550/arxiv.2512.22399

Wei, W., & Liu, L. (2024). Trustworthy distributed AI systems: Robustness, privacy, and governance. ACM Computing Surveys, 57(6), 1. https://doi.org/10.1145/3645102

Weiss, G. (Ed.). (1999). Multiagent systems: A modern approach to distributed artificial intelligence. MIT Press.

Winfield, A., & Nembrini, J. (2006). Safety in numbers: Fault-tolerance in robot swarms. International Journal of Modelling Identification and Control, 1(1), 30. https://doi.org/10.1504/ijmic.2006.008645

Wingerden, E. van, & Vigneswaran, D. (2024). The terrestrial trap: International relations beyond Earth. Review of International Studies, 50(3), 600. https://doi.org/10.1017/s0260210524000184

Wu, J., & Durfee, E. H. (2010). Resource-driven mission-phasing techniques for constrained agents in stochastic environments. Journal of Artificial Intelligence Research, 38, 415. https://doi.org/10.1613/jair.3004

Wyder, P. M., Bakhda, R., Zhao, M., et al. (2024). Robot metabolism: Towards machines that can grow by consuming other machines. arXiv. https://doi.org/10.48550/arxiv.2411.11192

Yap, X.-S., & Truffer, B. (2022). Contouring 'earth-space sustainability.' Environmental Innovation and Societal Transitions, 44, 185. https://doi.org/10.1016/j.eist.2022.06.004

Zambonelli, F., Omicini, A., Anzengruber, B., et al. (2014). Developing pervasive multi-agent systems with nature-inspired coordination. Pervasive and Mobile Computing, 17, 236. https://doi.org/10.1016/j.pmcj.2014.12.002

— End of Manuscript —