

PhD Research Proposal

Multi-Era Architectural Intelligence for AI-Driven Procedural Worldbuilding in Video Games

Candidate	[Your Name]
Supervisors	[Primary Supervisor], [Co-Supervisor]
Department	Computer Science / Digital Arts & Sciences
Institution	[University Name]
Program	Doctor of Philosophy (PhD)
Projected Duration	4–8+ years Core (Structural Inference): Y1–Y3 Style Applications: Y3–Y4 Integration & Demo: Y4+ Thesis: Y4–Y8+
Version	Final Draft, 2025

Abstract

Procedural Content Generation (PCG) has transformed game development by enabling vast environments at reduced production cost. Yet despite decades of research, existing systems remain architecturally shallow: they produce plausible surfaces and terrain while leaving the deeper structural logic of buildings (load-bearing systems, circulation hierarchies, spatial zoning) entirely unaddressed. This PhD proposes a framework for Multi-Era Architectural Intelligence (MEAI) built on a single, hard, foundational capability: the ability of an AI system to infer a building's internal structural skeleton (pillars, cores, load paths, circulation) from external geospatial data alone.

This structural inference capability is the research's primary and most demanding contribution, grounded in Earth's modern urban fabric as the base training and evaluation domain. It is anticipated to occupy the first three years of the program. Once established, the structural skeleton becomes a shared substrate onto which two concurrent style-application layers are built: era-shifting morphological transformation (reinterpreting the skeleton across historical and speculative periods) and xeno-biological adaptation (reconfiguring the skeleton for non-human species requirements). Both are treated as stylistic and parametric applications of the core inference engine, developed concurrently in Years 3–4. Planetary-scale generation is the integration demonstration: entire alien worlds whose architectural ecosystems are generated by applying era and xeno layers on top of inferred structural grammars. Thesis writing begins from Year 2 and intensifies from Year 4, with submission anticipated between Years 4 and 8 depending on the depth of industry integration pursued.

1. Introduction

Procedural Content Generation has become indispensable in commercial game development. Titles such as Minecraft, No Man's Sky, and Dwarf Fortress demonstrate that procedural systems can generate worlds of extraordinary scale and variety. However, a fundamental limitation persists: the architectural fabric of those worlds (the buildings, megastructures, and habitats that characters inhabit) lacks structural intelligence. Procedural buildings are essentially decorated boxes, visually convincing but structurally hollow.

Real architecture is not merely a surface phenomenon. It is governed by load distribution, material logic, circulation flow, spatial hierarchy, and cultural program. When these principles are absent, generated worlds feel inauthentic in ways that players perceive even if they cannot articulate them. A collapsed ruin that defies gravity, an alien settlement that is merely a human city with a texture swap, a planet whose cities are indistinguishable from one another: these are symptoms of the same underlying problem.

This research proposes a new paradigm: AI-Driven Architectural Intelligence, grounded in a single foundational capability: structural inference from geospatial data. The system learns to read a city's external observable data and predict the structural grammar governing its buildings: load paths, circulation cores, spatial hierarchies, zoning logic. Earth's modern urban fabric serves as the base training domain, providing the richest available combination of geospatial data, BIM records, and ground-truth architectural documentation. Once that structural intelligence is established, era-shifting and xeno-biological adaptation become applications layered on top, stylistic and parametric reinterpretations of the same underlying skeleton, applied to other historical periods, speculative futures, and alien worlds. Planetary generation is the integration demonstration: the proof that the framework scales to civilizational scope. This is a multi-year, multi-phase program of commensurate ambition.

2. Problem Statement

Contemporary PCG systems fail along four interconnected axes:

- RQ1. **Structural Shallowness.** Systems generate exterior forms and surface decoration but cannot infer load-bearing skeletons, structural grids, or circulation cores from building footprints or imagery. The result is architecture that looks plausible from a distance but collapses under any structural scrutiny.
- RQ2. **Era-Inconsistent Transformation.** Style transfer and reskinning pipelines change visual appearance without touching architectural logic. A medieval reskin of a modern tower block produces a building with anachronistic structural behavior: pointed arches on a steel-frame plan, or Roman columns supporting a curtain-wall grid.
- RQ3. **Human-Centric Bias.** All major PCG frameworks assume human ergonomic parameters: ceiling heights of 2.4–3.0 m, bilateral symmetry, gravity-oriented circulation. This makes them unfit for non-human civilizations, whether historical (Maya, Minoan), fictional (Halo Forerunners, Mass Effect Protheans), or speculative alien species.
- RQ4. **Poor Planetary Scalability.** Existing systems cannot generate a coherent architectural language across an entire planet, one that respects regional variation in climate, species physiology, social hierarchy, and historical development while maintaining civilizational coherence.

Together these limitations prevent game studios from realising genuinely believable large-scale worlds, and prevent researchers from studying architecture as a generative, reasoning-driven system rather than a decorative output.

3. Research Questions

The research is organized around five primary questions, sequenced to reflect the phased methodology described in Section 7:

- RQ5. **Structural Inference.** How can an AI model reliably infer the internal structural logic of a building (load-bearing elements, circulation cores, structural grid) from external geospatial maps, footprint data, and street-level imagery?
- RQ6. **Era-Consistent Transformation.** How can architectural logic be preserved, adapted, or plausibly degraded when a structure is transformed across historical, speculative, or alien eras, such that the result respects the material constraints and design norms of the target era?
- RQ7. **Xeno-Biological Adaptation.** How can PCG systems parameterize and adapt spatial environments to the biological and cultural requirements of non-human species, including non-standard ergonomics, atmospheric conditions, mobility patterns, and social hierarchies?
- RQ8. **Planetary-Scale Coherence.** How can species physiology, biome variation, social organization, and civilizational history be encoded into a generative grammar capable of producing coherent architectural ecosystems at planetary scale?
- RQ9. **Game-Ready Integration.** How can the above capabilities be packaged into real-time or near-real-time tools that integrate with existing game engines and development pipelines, enabling new gameplay mechanics and reducing production cost?

4. Research Vision: Multi-Era Architectural Intelligence

The overarching vision is to shift PCG from surface-level decoration to deep architectural reasoning. The framework is organized as a layered pyramid: structural inference at the base, style applications (era and xeno) in the middle, and planetary-scale generation at the apex. Each layer depends on and extends the one beneath it. The base layer is grounded in the most data-rich, verifiable domain available: Earth's modern urban environment. Everything above it (historical eras, alien species, other planets) is a structured transformation or parameterization of that base.

4.1 Foundation Layer: Structural Inference from Geospatial Data (Earth Modern Cities)

This is the primary and most demanding contribution of the research. The foundation layer focuses exclusively on Earth's contemporary built environment: the modern city. This choice is deliberate. Modern cities offer the densest available combination of satellite imagery, OpenStreetMap footprints, CityGML semantic models, LiDAR point clouds, BIM repositories, and ground-truth structural documentation. They are the ideal training domain for a system that needs to learn the relationship between external observable form and internal structural logic.

The AI infers from external geospatial data:

- Load-bearing pillar and column locations
- Core wall and shear wall placement
- Structural grid spacing and orientation

- Lift lobbies, stairwells, and circulation cores
- Functional zoning and room adjacency patterns
- Floor-plate typology (perimeter-core, open plan, cellular)

Validation draws on datasets where both external and internal documentation are available: BuildingNet, CityGML models, and BIM repositories. The output is a structural skeleton graph, the shared lingua franca by all upstream modules. This base layer is expected to occupy Years 1–3 of the program. It is genuinely hard: the inference problem is under-constrained, the ground truth is expensive to obtain, and the evaluation methodology (what does it mean for an inferred skeleton to be ‘correct’?) must itself be developed as part of the research contribution.

4.2 Style Application Layer A: Era-Shifting Transformation

Once a structural skeleton exists, era-shifting is a stylistic reinterpretation: the same underlying load logic expressed through different material vocabularies, construction technologies, and design norms. This layer is the primary mechanism for generating other-planet and historical settings in the framework: an alien world is, at its structural core, a skeleton derived from the same inference grammar as a modern city, then shifted into a speculative or historical style register. This keeps the hard problem of structural inference cleanly separated from the style problem of transformation.

Target transformation registers include:

- Prehistoric and vernacular (timber, earth, stone)
- Classical antiquity (column-and-lintel, vaulted masonry)
- Medieval (Gothic vaulting, load-bearing masonry)
- Industrial revolution (cast-iron and wrought-iron frames)
- Modernist (reinforced concrete, steel frame, curtain wall)
- Sci-fi near-future (advanced composites, active structural systems)
- Post-apocalyptic (structural degradation, improvised repair)
- Deep speculative (non-Euclidean, energy-field structures)

Transformations are governed by constraint satisfaction: material limits, gravity assumptions, era-appropriate connection logic. Degraded states (ruins, partial collapse) model progressive load-path failure. This layer is developed concurrently with 4.3 in Years 3–4.

4.3 Style Application Layer B: Xeno-Biological Adaptation (Concurrent with 4.2)

Xeno-biological adaptation is the second style-application layer, developed concurrently with era-shifting. It reconfigures the structural skeleton to accommodate non-human biological and cultural requirements. Where era-shifting changes the material and aesthetic register of a skeleton, xeno-biological adaptation changes its spatial geometry, proportion, and programmatic organization. Both layers operate on the same structural skeleton output from the foundation layer.

Species parameters encoding:

- Body plan: height, limb count, locomotive mode (bipedal, quadrupedal, arboreal, aquatic, aerial)

- Sensory profile: visual spectrum, echolocation, pressure sensitivity
- Atmospheric requirements: gas mix, pressure range, temperature envelope
- Social structure: hierarchy depth, privacy norms, ritual spatial requirements
- Cultural logic: orientation taboos, sacred geometries, threshold symbolism

Outputs are validated against biological plausibility (can this species inhabit this space?) and cultural coherence criteria. The concurrent development of 4.2 and 4.3 allows them to be combined: an alien world in a historical-equivalent era requires both transformations simultaneously.

4.4 Integration Demonstration: Planet-Scale Civilizational Ecosystems

Planetary-scale generation is the capstone demonstration of the framework rather than a research problem in its own right. Its purpose is to prove that structural inference, era transformation, and xeno-biological adaptation, applied at scale with a hierarchical generation architecture, produces coherent civilizational worlds. It runs concurrently with Phases 2 and 3 and forms the primary content of the thesis demonstration chapters.

An important clarification governs the scope of this layer. This research is concerned with the architecture of civilization: the man-made structures, settlements, and built environments that intelligent species construct. It is not a framework for generating planetary terrain, geology, weather systems, or natural environments in their own right. However, the natural character of a planet is treated as a key input parameter that shapes architectural output. On Earth, climate, terrain, available materials, and gravitational constant have always driven how humans build: heavy masonry in earthquake zones, steep-pitched roofs in high-snowfall regions, stilted structures in flood plains, passive cooling geometries in desert climates. The same logic applies at planetary scale. A high-gravity world produces squat, low-profile structures with heavy load transfer; a low-gravity world permits tall, slender forms with minimal structural redundancy; a methane-atmosphere planet drives sealed, pressurized envelopes regardless of species. The planet informs the building. The framework reads planetary parameters as constraints and generates architecturally authentic responses to them. Generating the planet itself (terrain, ecosystem, geology) is a separate and complementary problem addressed by existing PCG terrain systems, and is explicitly out of scope here.

- Biome and climate zones driving material choices and thermal design
- Historical development trajectories (founding, expansion, decline) applied via era layer
- Hierarchical settlement typology: outpost, town, city, megacity, arcology
- Civilizational aesthetic coherence maintained by shared structural grammar

5. Literature Review

The research sits at the intersection of five established fields. This review situates the proposed contribution relative to each and identifies the specific gaps it addresses.

5.1 Procedural Content Generation

The foundational PCG literature establishes grammar-based generation (Prusinkiewicz & Lindenmayer, 1990; Müller et al., 2006), noise-based terrain (Perlin, 1985), and agent-based simulation (Togelius et al., 2011). Shaker, Togelius & Nelson (2016) provide the definitive survey. Commercial applications include the L-system-derived buildings of CityEngine, the Whittaker biome model used in No Man's Sky, and the emergent narrative engine of Dwarf Fortress. The critical limitation identified across this body of work is consistent: PCG systems operate on surface appearance, not structural reasoning.

5.2 Urban and Architectural Generation

Parish & Müller (2001) established the road-network-then-parcel-subdivision pipeline that underlies most procedural city work. Wonka et al. (2003) introduced shape grammars for building facades. Merrell et al. (2010) extended this to 3D interior layout. More recently, Liu et al. (2022) demonstrated graph neural networks for floor plan generation, and Nauata et al. (2021) applied transformer architectures to room layout prediction. None of these systems reason about structural load paths or cross-era transformation.

5.3 AI-Driven Generation

Deep learning has dramatically expanded PCG capability. Generative Adversarial Networks (Goodfellow et al., 2014) and diffusion models (Ho et al., 2020) enable high-fidelity 2D and 3D synthesis. Reinforcement learning has been applied to level generation (Khalifa et al., 2020). Large language models (Brown et al., 2020) show promise for rule generation and narrative scaffolding. Graph Neural Networks (Kipf & Welling, 2017) are the natural architecture for structural skeleton reasoning. The proposed research integrates these advances into a coherent architectural reasoning framework, an integration that has not previously been attempted.

5.4 Style Transfer and Temporal Transformation

Neural style transfer (Gatys et al., 2015) provides compelling visual-aesthetic transformation but operates on pixel statistics, not structural grammar. Procedural style grammars (Stiny, 1980) offer rule-based transformation but require manual rule specification. Architectural morphology research (Hillier & Hanson, 1984; Steadman, 2014) provides theoretical frameworks for understanding building typology evolution. The proposed research synthesizes these traditions into a learned, constraint-aware transformation system.

5.5 Biological and Cultural Adaptation

Anthropometric design (Pheasant & Haslegrave, 2006) establishes the human ergonomic baseline. Environmental psychology (Gifford, 2014) models how spatial configuration affects occupant behavior. Speculative xenobiology (Cohen & Stewart, 2002) provides frameworks

for reasoning about non-terrestrial biologies. Space syntax (Hillier, 1996) offers quantitative tools for analysing spatial hierarchy and circulation. These fields collectively provide the theoretical grounding for the xeno-biological adaptation module.

5.6 Research Gap

No existing PCG framework integrates structural inference, era-consistent transformation, non-human biological adaptation, and planetary-scale generation into a unified system. This proposal addresses that gap directly. The nearest adjacent work, specifically Liu et al. (2022) on graph-based floor plans and Ritchie et al. (2019) on indoor scene synthesis, addresses isolated sub-problems within a human-centric, single-era, building-scale frame.

6. Methodology

The methodology reflects the layered architecture of the research: a deep foundation phase, two concurrent style-application phases, and an integration demonstration that runs in parallel from Year 3 onward. Thesis writing is treated as an ongoing activity beginning in Year 2, not a terminal phase. Each stage produces both a publishable output and a software module contributing to the MEAI framework.

6.1 Stage 1: Foundation: Structural Inference from Earth Modern Cities (Years 1–3)

This is the research's core contribution and its longest stage. The domain is deliberately constrained to Earth's contemporary urban fabric. Modern cities are chosen because they offer the most tractable ground truth: satellite imagery, OpenStreetMap footprints (550M+ buildings globally), CityGML semantic models, LiDAR surveys, and BIM repositories can be cross-referenced to construct training labels for structural elements. This domain constraint is a research design choice, not a limitation; it enables rigorous validation before the framework is extended to speculative settings.

Stage 1 work includes:

- Dataset construction: aggregate BuildingNet, OpenStreetMap, CityGML (Berlin, New York, Singapore, Helsinki), LiDAR corpora, and a curated BIM dataset. Develop annotation protocols for ground-truth structural labels in collaboration with structural engineering partners.
- Evaluation methodology development: define what 'correct' structural inference means. This is itself a novel contribution, as there is no accepted benchmark for this task. Proposed metrics include load-path continuity, grid regularity, structural engineer agreement rate, and FEA proxy scores.
- Model development: train a Graph Neural Network (or successor architecture) over building footprint graphs to predict structural skeleton labels: pillar, core wall, grid node, circulation node, slab zone. Benchmark against rule-based baselines and human expert annotations. Target: structural inference agreement rate > 0.80 with qualified structural engineers on held-out city regions.
- Iterative refinement: the three-year window allows multiple training cycles, architecture experiments, and dataset expansions. Year 1 targets proof-of-concept on a single city typology (e.g. North American grid-plan); Year 2 extends to diverse global typologies; Year 3 hardens the model and prepares the publication.

Deliverable: StructureNet (trained model, evaluation benchmark, dataset) + journal paper (IEEE TVCG or ISPRS Journal of Photogrammetry and Remote Sensing). Thesis chapters 1–3 drafted during this stage.

6.2 Stage 2A: Style Layer: Era-Shifting Transformation (Years 3–4, concurrent with 2B)

With a structural skeleton as input, era-shifting is reframed as a conditional style generation problem: given a skeleton and an era target, produce a geometrically and materially plausible realization. The key insight is that the skeleton does not change between eras; only its materialization does. This dramatically reduces the problem complexity relative to systems that attempt full structural re-generation.

Stage 2A work includes:

- Era style dataset: pair historical architectural surveys, archaeology databases, and architectural history corpora with modern structural equivalents to create era-labeled training pairs.
- Constraint layer: encode material limits, gravity assumptions, and connection logics per era as hard constraints on the generative model.
- Model: conditional diffusion model or style-conditioned VAE taking (skeleton, era) → 3D architectural realization. Evaluate on structural plausibility (FEA proxy) and era authenticity (architectural historian review panel).

Deliverable: EraShift module + conference paper (ACM SIGGRAPH Asia or Eurographics).

6.3 Stage 2B: Style Layer: Xeno-Biological Adaptation (Years 3–4, concurrent with 2A)

Xeno-biological adaptation applies the same architectural skeleton to a non-human species parameter profile. It is developed concurrently with Stage 2A and shares its evaluation timeline. The two style layers are designed to be composable: a single structure can simultaneously receive an era transformation and a species adaptation, producing, for example, a Forerunner-era high-gravity-adapted megastructure.

Stage 2B work includes:

- Species parameter schema: developed in collaboration with game designers and speculative biologists. Grounded in non-human Earth architectural precedents (termite mounds, coral reef structures, animal burrow networks) as proxy validation cases.
- Model: species-conditioned architectural grammar mapping (skeleton, species profile) → spatially adapted realization. Evaluate via biological plausibility scoring and user studies with game designers.
- Composition test: validate that EraShift and XenoArch can be applied simultaneously without conflict, producing coherent compound outputs.

Deliverable: XenoArch module + conference paper + design tool prototype (FDG or IEEE CIG).

6.4 Stage 3: Integration Demo: Planetary-Scale Generation (Years 4+, concurrent)

Planetary generation is the integration demonstration. It does not introduce new AI research problems; it applies the foundation and style layers at scale, with a hierarchical generation architecture managing consistency across settlement types, biomes, and civilizational histories. The primary research contribution here is the pipeline architecture and the evaluation framework for world-scale coherence.

- Hierarchical PCG pipeline: settlement networks, biome-responsive style selection, civilizational history simulation driving era-layer inputs.
- Coherence evaluation: proposed novel metric suite for civilizational architectural consistency (structural grammar continuity, style drift plausibility, biological consistency across settlement types).

- Game engine integration: Unreal Engine 5 and Unity plugins wrapping all three modules. Structure Scanner, Era Shift Tool, and Species Configurator as gameplay-facing tools.

Deliverable: PlanetGen pipeline + game prototype + open-source SDK. Primary thesis demonstration content.

6.5 Thesis Writing Strategy

Thesis writing is not a terminal activity. The following schedule ensures the thesis is substantially drafted before the integration phase completes:

- Year 2: Literature review and problem statement chapters drafted and reviewed by supervisors. StructureNet evaluation methodology chapter written in parallel with model development.
- Year 3: StructureNet results chapter completed. Stage 2A/B methodology chapters drafted concurrently with experimental work.
- Year 4: Style layer results chapters and integration architecture chapter. Minimum viable thesis is defensible at this point for candidates on a 4-year program.
- Years 5–8+: Extended thesis incorporating full planetary demonstration, industry evaluation, and deeper theoretical synthesis. Candidates pursuing deeper industry integration or additional publications may submit in this window.

7. Research Timeline

The timeline reflects the layered structure of the research. Stage 1 (structural inference) is the longest and most demanding. Stages 2A and 2B run concurrently. Stage 3 (integration demo) begins concurrent with Stage 2 and continues to thesis submission. Thesis writing is an ongoing activity throughout.

Phase	Duration	Year(s)	Key Milestones
Stage 1: Structural Inference (Earth modern cities, foundation)	36 months	Y1–Y3	Dataset + benchmark built; StructureNet trained; evaluation methodology published; TVCG paper submitted
Stage 2A: Era-Shifting (style layer, concurrent)	18 months	Y3–Y4	EraShift module; constraint solver; era-annotated dataset; SIGGRAPH Asia paper
Stage 2B: Xeno-Bio Adaptation (style layer, concurrent with 2A)	18 months	Y3–Y4	Species schema; XenoArch module; design tool prototype; user study; FDG paper
Stage 3: Planetary Demo (integration, runs concurrent from Y3)	24–36 months	Y4–Y6+	PlanetGen pipeline; coherence metrics; engine plugins; game prototype; SDK release

Thesis Writing (ongoing, not terminal)	Continuous	Y2–Y8+	Lit review + problem statement (Y2); StructureNet chapters (Y3); style layer chapters (Y4); integration + synthesis (Y4–Y8+)
---	------------	--------	--

The minimum viable PhD is defensible at the end of Year 4: StructureNet (foundation) + EraShift + XenoArch (style layers) constitute three novel contributions with published papers. The extended thesis (Years 5–8+) adds the planetary integration demonstration, deeper industry evaluation, and broader theoretical synthesis. The 4–8+ year window is an honest reflection of frontier AI research uncertainty, not schedule padding.

8. Expected Contributions

8.1 Academic Contributions

- StructureNet: a graph neural network architecture and benchmark for structural skeleton inference from geospatial data (the primary novel contribution, grounded in Earth modern cities)
- EraShift: a constraint-aware conditional generative model for era-consistent architectural style transformation, applied to non-Earth and historical settings
- XenoArch: a species-parameterized architectural grammar for non-human spatial generation, composable with EraShift
- PlanetGen: a hierarchical PCG integration pipeline demonstrating civilizational-scale coherent generation
- A unified MEAI theoretical framework articulating the pyramid architecture: structural inference → style application → planetary demonstration

8.2 Industry Contributions

- Open-source MEAI SDK compatible with Unreal Engine 5 and Unity
- Structure Scanner, Era Shift Tool, and Species Configurator as game-ready tools
- Documentation and case studies for integration into AAA and indie production pipelines
- Demonstrated relevance to studios including 343 Industries, Bethesda, Rockstar Games, CD Projekt Red, and miHoYo

9. Game Concepts Enabled by the Framework

This section demonstrates the research's applied significance. Each concept below is enabled by one or more MEAI modules and represents a genre that is currently commercially infeasible without this research.

- **Multi-Era Structural Reconstruction:** Players explore a single building across five or more historical eras, with AI reconstructing structurally coherent versions of each. Gameplay rewards structural literacy: understanding load paths to navigate ruins, identify weaknesses, or reconstruct collapsed sections.
- **Xeno-Biological Civilization Builder:** A strategy game in which players design cities for procedurally generated alien species. The AI enforces biological plausibility: a methane-breathing species cannot have open courtyards on an oxygen world, creating emergent constraint-driven design challenges.
- **Planet-Scale Exploration RPG:** An open-galaxy RPG in which every planet has a coherent civilizational history expressed through its architecture. Ruins indicate prior inhabitants; settlement patterns encode social hierarchies; material choices reflect biome constraints.
- **Archaeological Investigation Game:** Players uncover and reconstruct alien ruins using structural inference tools. The AI generates ruins as degraded structural skeletons; players reason about original forms from surviving load-bearing elements.

- **Adaptive Sci-Fi Shooter:** Combat arenas that dynamically adapt their spatial configuration to the species currently occupying them, rewarding players who understand alien ergonomics and architectural logic.

10. Industry Case Study: Halo: Reclaimer's Odyssey

This motivating case study illustrates how the MEAI framework could power a next-generation AAA title. It is presented as a design vision, not a commercial commitment, to demonstrate the practical stakes of the research and its alignment with current industry directions.

10.1 Concept Overview

Halo: Reclaimer's Odyssey is a planet-scale open-galaxy RPG built on the MEAI framework. Its central design tension is structured around three civilizational forces, each with a distinct architectural intelligence; and within the third force, a rich coalition of species whose built environments are as biologically and culturally distinct as the species themselves:

- **The Atom (Forerunner):** Hard-light, modular, machine logic. Monumental scale with non-Euclidean circulation, distributed structural redundancy, and recursive geometric self-similarity. Forerunner architecture is generated using the deep-speculative era module: structures that predate human civilization by 100,000 years, built to last geological timescales, with spatial hierarchies organized around the logic of a machine god rather than biological need.
- **The Cell (Flood):** Architecture as living tissue. The Flood does not build; it consumes and mutates. Load paths dissolve into mycelial growth networks; spatial hierarchy collapses into undifferentiated hive pressure; walls become membranes. Generated using the xeno-biological adaptation module with Flood bio-mass parameters: zero ergonomic constraints, infinite redundancy, environments defined entirely by infection vectors and biomass distribution rather than circulation or function.

The third force is not a monolith but a coalition of distinct civilizations, each contributing a unique architectural identity to the galaxy:

- **Humanity / UNSC:** Industrial-functional architecture grounded in Earth modern typology (the base domain of the StructureNet module. UNSC settlements are the framework's structural inference anchor: grid-plan cities, reinforced-concrete military installations, modular orbital stations. Recognizable, load-bearing, pragmatic. The structural skeleton that all other civilizations are measured against.
- **San'Shyuum (Prophets):** Theocratic and ceremonial. The ruling caste of the former Covenant Empire designed their spaces around religious authority and political control: high vaulted processional corridors, tiered audience chambers, sacred geometries derived from Forerunner iconography. San'Shyuum architecture is the era-shifting module's key test case: a civilization that copied Forerunner aesthetics without understanding Forerunner structural logic, producing structurally imitative but functionally distinct spaces.
- **Sangheili (Elites):** Ceremonial verticality and martial symmetry. Sangheili society is organized around honor, lineage, and ritualized combat, and their architecture encodes all three: tall vertical circulation with open sight-lines for combat readiness, bilateral symmetry reflecting rank hierarchy, threshold spaces (gates, antechambers) of disproportionate scale that encode the significance of crossing them. Keep-cities on Sanghelios are the model: fortified, vertical, with communal great halls elevated above private warrens.

- **Jiralhanae (Brutes):** Brutalist and tribal-aggressive. Jiralhanae architecture reflects a species organized around pack dominance hierarchies: low ceilings in subordinate spaces, oversized chieftain chambers, no distinction between armory and living space. Settlements are improvised, heavily fortified, with structural redundancy driven by anticipation of internal violence rather than external threat. Generated using the xeno-biological adaptation module with Jiralhanae physiology parameters: large bipedal mass, high aggression coefficient, short planning horizons.
- **Unggoy (Grunts):** Methane warrens and pressure-sealed habitats. Unggoy are methane-breathing bipeds from the high-pressure world of Balaho. Their architecture is defined entirely by atmospheric containment: low ceilings scaled to 1.4m average height, pressure-lock threshold sequences at every transition, communal sleeping warrens with no private space, methane recycling infrastructure embedded in every wall. The densest, most biologically constrained architectural profile in the framework.
- **Kig-Yar (Jackals):** Avian-pirate pragmatism. Kig-Yar are opportunists, and their settlements reflect it. Lightweight, modular, reconfigurable structures optimized for rapid assembly and abandonment. Elevated walkways and perch-height platforms reflecting avian postural preferences. No monumental scale, no ceremonial space; everything is storage, surveillance, or escape route. The xeno-biological module must account for digitigrade limb structure and acute visual acuity driving spatial sightline optimization.
- **Mgalekgolo / Lekgolo (Hunters):** Gestalt colonial architecture. Lekgolo are colonial worm organisms that form massive gestalt bodies (Mgalekgolo) for combat and labor. Their spatial needs are radically unlike any other species: no individual-scale rooms, no human-proportioned corridors. Lekgolo settlements (rare and poorly documented) are network structures of tunnels and chambers sized for colony-mass movement, with no concept of private space. The most structurally novel profile in the framework.
- **Yanme'e (Drones):** Aerial insectoid hive structures. Yanme'e are flying insectoids that originally had an interplanetary civilization before Covenant incorporation. Their architecture is three-dimensional in a way no ground-bound species produces: no floor-plan logic, vertical and horizontal surfaces treated identically, spaces organized around queen-chamber hierarchies with radial circulation patterns. Gravity is a constraint to be minimized, not designed around.
- **Huragok (Engineers):** Forerunner-created, non-combatant, floating. Huragok were created by the Forerunners as maintenance organisms and have no homeworld architecture of their own; they exist within the architecture of others, maintaining it. Their presence in the framework is as an adaptive layer: spaces that have been occupied and maintained by Huragok show distinctive modification patterns: access panels, maintenance channels, reconfigured load paths. The xeno-biological module uses Huragok occupation signatures as a temporal marker of Forerunner-era habitation.

Scope Note: Spacecraft and Vehicle Design. *Each civilization and species in this framework would logically extend its architectural identity into its spacecraft, vehicles, and mobile habitats. Forerunner keyships, Covenant cruisers, UNSC frigates, Kig-Yar raider craft, and Flood-infected vessels each carry the same biological and cultural logic as their ground-based architecture: the same structural grammar, spatial proportions, and material vocabulary applied to a mobile platform. Generating species-authentic spacecraft design using the same parametric framework is a compelling and natural extension of this research. However, vehicle and spacecraft design constitutes a distinct domain with its own engineering constraints (propulsion, atmospheric re-entry, zero-gravity spatial logic, hull stress modeling) that falls outside the scope of this PhD. It is noted here as a clear direction for future research and as a commercially significant adjacent problem for game studios.*

Scope Note: Planetary Nature and Natural Environments. *This research is fundamentally about the architecture of civilization: the man-made structures, settlements, and built environments that intelligent species construct. It is not a framework for generating planetary terrain, geology, weather systems, natural landscapes, or ecosystems. Those domains are well-served by existing PCG terrain systems and are explicitly out of scope here. However, the natural character of a planet is a primary input parameter that shapes all architectural output. On Earth, building design has always been a response to nature: heavy masonry and low profiles in earthquake-prone regions, steep-pitched roofs under high snowfall, stilted structures in flood plains, passive cooling geometries and thick thermal mass in desert climates, elevated foundations in permafrost. The same relationship holds at planetary scale. A high-gravity world produces squat, structurally redundant forms; a low-gravity world permits tall, slender structures with minimal load transfer; a methane atmosphere drives sealed pressurized envelopes regardless of species; a tidally locked planet creates permanent day-side and night-side architectural vocabularies. The planet defines the constraints. The MEAI framework reads those constraints as parameters and generates architecturally authentic responses to them. What it does not do is generate the planet itself. In a complete game pipeline, a terrain and environment system (existing technology) would feed planetary parameters into the MEAI framework, which would then generate the civilization built on top of that world.*

The case study maps each research module to a specific game system:

- **StructureNet:** Powers the in-game Structure Scanner tool, allowing players to detect load paths, identify maintenance shafts, and discover hidden vaults in Forerunner facilities.
- **EraShift:** Generates the temporal layering of each world: Forerunner construction era, Covenant occupation era, post-Covenant collapse, and current state, all structurally consistent with their respective material constraints.
- **XenoArch:** Ensures every settlement type is biologically authentic across the full species roster: Sangheili keep-cities prioritize vertical ceremonial space and bilateral martial symmetry; Unggoy warrens are methane-sealed, low-ceilinged, and communal; Jiralhanae outposts reflect pack dominance hierarchies in their spatial proportions; Kig-Yar platforms optimize for avian perch-height and sightlines; Yanme'e hives ignore floor-plan logic entirely; Lekgolo tunnel networks are sized for colony-mass movement; Flood hives dissolve all conventional spatial hierarchy into biomass pressure.
- **PlanetGen:** Generates the full galaxy, spanning thousands of star systems, each with coherent planetary civilizational histories, at a scale and consistency impossible with hand-crafted content.

10.3 Industry Alignment

This case study demonstrates alignment with the production needs of major studios currently developing open-world titles at scale. The structural inference and era transformation tools are directly applicable to Bethesda's open-world pipelines, Rockstar's city generation systems, CD Projekt Red's environmental storytelling workflows, and miHoYo's multi-civilizational world design. The academic research thus has a clear and immediate path to industry impact.

11. Risk Analysis and Mitigation

Risk	Likelihood / Impact	Mitigation
Insufficient labeled structural data for Stage 1 training	Medium / High	Develop synthetic data pipeline from CityGML + BIM; crowdsource annotation via structural engineering schools; use self-supervised pre-training on unlabeled footprints
Evaluation methodology for structural inference cannot be agreed by experts	Medium / High	This is itself a research contribution; establish multi-criteria protocol early (Y1) with structural engineering advisory panel
Era/xeno style layers degrade structural plausibility of base skeleton	Medium / Medium	Enforce skeleton as hard constraint in generative model; era and xeno layers operate on materialization only, not topology
Concurrent Stage 2A/2B development creates integration conflicts	Low / Medium	Define shared skeleton schema at start of Stage 1; design composition API before Stage 2 begins
Real-time performance insufficient for game integration	High / Medium	Decouple offline generation (inference, transformation) from runtime lookup; use LOD streaming and precomputed skeleton libraries
Thesis scope expands beyond what is defensible at Y4 minimum	Medium / High	Define Y4 minimum viable thesis explicitly in Year 1 plan; annual supervisory scope review; planetary demo treated as extension not requirement

12. Resources and Infrastructure

12.1 Computational

- GPU cluster access: minimum 4x NVIDIA A100 or equivalent for training phases
- Cloud compute budget: estimated £40,000–60,000 over project duration
- Storage: 50 TB minimum for geospatial and 3D architectural datasets

12.2 Datasets

- BuildingNet (3D building mesh dataset with semantic labels)
- OpenStreetMap global footprint corpus (550M+ buildings)
- CityGML semantic 3D city models (Berlin, New York, Helsinki, Singapore)

- ShapeNet Buildings subset
- Proprietary BIM data (to be negotiated with architectural firms)
- Historical architectural surveys (archive digitization program proposed)

12.3 Collaborators

- Architecture department: structural logic validation and historical era annotation
- Game industry partners: user studies, production pipeline integration, commercialisation
- Speculative biology consultants: xeno-biological parameter validation

13. Conclusion

This proposal is built around a single, honest conviction: structural inference is the hard problem, and everything else follows from it. By grounding the research in Earth's modern urban fabric (the richest, most verifiable architectural dataset available) and spending three years building a system that can genuinely reason about structural logic from external data alone, this program creates a foundation that justifies every subsequent application. Era-shifting and xeno-biological adaptation are not independent research problems bolted onto the side: they are stylistic reinterpretations of structural skeletons, and their development becomes tractable precisely because the foundation is solid.

The 4–8+ year timeline is not ambition inflation. Structural inference from geospatial data is an under-constrained, under-benchmarked problem with no accepted evaluation protocol. Solving it rigorously, publishing it, and building the two style-application layers on top constitutes a minimum viable PhD in three to four years. The planetary integration demonstration and thesis synthesis (the full vision of the research) may take eight. Both outcomes are intellectually defensible. Both would be meaningful contributions to AI, PCG research, and the game industry.

Procedural Content Generation has given the game industry scale. Multi-Era Architectural Intelligence proposes to give it depth, one structural skeleton at a time.

References

- Brown, T. et al. (2020). Language Models are Few-Shot Learners. NeurIPS 33.
- Cohen, J. & Stewart, I. (2002). *Evolving the Alien*. Ebury Press.
- Gatys, L., Ecker, A. & Bethge, M. (2015). A Neural Algorithm of Artistic Style. arXiv:1508.06576.
- Gifford, R. (2014). *Environmental Psychology: Principles and Practice*. Optimal Books.
- Goodfellow, I. et al. (2014). Generative Adversarial Networks. NeurIPS 27.
- Hillier, B. (1996). *Space is the Machine*. Cambridge University Press.
- Hillier, B. & Hanson, J. (1984). *The Social Logic of Space*. Cambridge University Press.
- Ho, J. et al. (2020). Denoising Diffusion Probabilistic Models. NeurIPS 33.
- Khalifa, A. et al. (2020). Procedural Content Generation via Reinforcement Learning. AIIDE.
- Kipf, T. & Welling, M. (2017). Semi-Supervised Classification with Graph Convolutional Networks. ICLR.
- Liu, C. et al. (2022). Towards Complete Scene Generation via Graph Neural Network. CVPR.
- Merrell, P. et al. (2010). Interactive Furniture Layout Using Interior Design Guidelines. SIGGRAPH.
- Müller, P. et al. (2006). Procedural Modeling of Buildings. SIGGRAPH.
- Nauata, N. et al. (2021). House-GAN++: Generative Adversarial Layout Refinement. CVPR.
- Parish, Y. & Müller, P. (2001). Procedural Modeling of Cities. SIGGRAPH.
- Perlin, K. (1985). *An Image Synthesizer*. SIGGRAPH.
- Pheasant, S. & Haslegrave, C. (2006). *Bodyspace: Anthropometry, Ergonomics and the Design of Work*. Taylor & Francis.
- Prusinkiewicz, P. & Lindenmayer, A. (1990). *The Algorithmic Beauty of Plants*. Springer.
- Ritchie, D. et al. (2019). Fast and Flexible Indoor Scene Synthesis via Deep Convolutional Generative Models. CVPR.
- Shaker, N., Togelius, J. & Nelson, M. (2016). *Procedural Content Generation in Games*. Springer.
- Steadman, P. (2014). *The Evolution of Designs*. Routledge.
- Stiny, G. (1980). *Introduction to Shape and Shape Grammars*. Environment and Planning B.
- Togelius, J. et al. (2011). Search-Based Procedural Content Generation. EuroGP.
- Wonka, P. et al. (2003). *Instant Architecture*. SIGGRAPH.