Next-Gen Landscape Design: Agentic AI + Digital Twins

Stephen Ervin[1]

[1]Harvard Graduate School of Design (Retired, Associate)/USA · stephen.ervin@gmail.com

Abstract: The complexity of issues confronting landscape architects and other environmental designers, in the face of the global polycrisisvare immense and growing. Artificial Intelligence (AI) tools, may represent a panacea, but they are still very little beyond infancy; results across many domains are so far quite mixed; and limitations and concerns are abundant. In this context, several different but related emergent technological advances do seem likely to be key drivers of the next generation of AI-enabled computer-aided-design tools intended to help design thinking and manage real-world complexity: 'Agentic AI' (AAI) platforms connected to 'Digital Twins', esp. 'Digital Twin Earths' (DTEs), with 'Model Context Protocol' (MCP) and 'Retrieval-Augmented Generation' (RAG) used to control and constrain the AI agents. Following a brief description of these technologies and their foundations, I explore some application possibilities by re-visiting a decade-old 'System for Geodesign' with AAI, MCP, RAG, and DTE integrated, and outline an updated prescription for a technologically enhanced and enabled next-gen 'AAI+DTE' design system (not just for geodesign). A highly simplified prototype agentic system implementation is described.

Keywords: Next-Gen Computer-Aided Landscape Design, Agentic AI, Digital Twin Earth, Retrieval-Augmented Generation, Model Control Protocol

1 Introduction

As the big-data, scale-hopping, information-processing, multitudinous detail-managing, and time-consuming collaboration-choreographing demands on digital landscape architects continue to grow, especially in the face of global environmental issues, the promise of computer-aided design to alleviate some time-consuming and repetitive tasks so as to give designers time and space to design is evermore seductive. From Sutherland's seminal SketchPad, through 2D and 3D CAD, GIS, Remote-Sensing and Image-Processing, GPS and then BIM, and now AI-in-CAD, some productivity gains have been achieved. Over the same time span, the complexity of issues confronting environmental designers, and the complexities of marshaling design and construction projects in the real world, have been even further increased. The current explosion of AI, with its similar promises, is still in its infancy; results are so far mixed; concerns are real; and speculation is rampant.

For the past few years, it has been apparent that the long-awaited potential of Artificial Intelligence (AI) as foreseen by (TURING 1948), (MINSKY 1986), and (SIMON 1996) may have begun to be delivered (albeit with some already apparent downsides!). The explosion of Machine Learning (ML), Large Language Models (LLMs), Generative Pre-Trained Transformers (GPTs), Generative Adversarial Networks (GANs), Diffusion Model Generative Pre-trained Autoregressive Diffusion Transformers (GPDiTs), Model Context Protocol (MCP), and Retrieval-Augmented Generation (RAG), along with other techniques and platforms including ChatGPT, Gemini, Claude, DALL-E, Stable Diffusion, and now a whole host of other technologies and acronyms, in just a few years' time, has absolutely rocked worlds – of de-

Journal of Digital Landscape Architecture, 11-2026, pp. 62-74. © Wichmann Verlag, VDE VERLAG GMBH · Berlin · Offenbach. ISBN 978-3-87907-770-0, ISSN 2367-4253, e-ISSN 2511-624X, doi:10.14627/537770006.

sign (CANTRELL et al. 2021), construction (ABIOYE et al. 2021), education (BOWEN & WATSON 2024), communication, entertainment, and everything else in between.

In recent years, a number of papers published in JoDLA and elsewhere have chronicled the emergence of AI in Landscape Architecture (BRAIDEN et al. 2025), (XING et al. 2025). This emergence is proceeding so rapidly that already these early reviews and experiments seem almost quaint.

What most AI at this moment in time (2025) is already very good at doing includes organizing, summarizing, modifying, and even synthesizing textual data and responding to natural language prompts, based on a huge library of background knowledge in ways which are usually coherent, well-structured, plausible, and believable, even if not always factually true. And, to a somewhat lesser extent, doing the same for images, music, computer code, and other media. Its facility with interaction and apparent understanding of the meaning of the text prompt (s) it is responding to are deceptively good, notwithstanding that we know, in fact, there is no real understanding or reasoning, only statistical similarity and likelihood, at work; and despite the many flaws and potential problems with the current technologies, which depend on a large body ('corpus') of training samples, possibly biased or incomplete, for all of their power.

As to whether and how the GenAI programs may be designing, there is much discussion; on this topic, ChatGPT-4.1 itself says: "[Generative AI]'s processes lack intentionality, contextual awareness, and critical reflection that are central to design as traditionally understood."

This observation is behind the oft-repeated suggestion that "there needs to be a human in the loop" (WIKIPEDIA 2025), and the assertion that "co-design between humans and machines is the logical evolution of computer-aided design". There is good evidence that already, AI tools *can* aid design and designers in various ways, and so belong in the modern Computer-Aided Design toolbox. How these tools will become sharper, more accurate, and effective, more widely and creatively used, etc., will doubtless continue to consume the Digital Landscape Architecture community, along with much of the rest of the world. The goal of the tools in the CAD toolbox is no longer just to "free designers from burdensome and repetitive tasks", and their promise is not to *reduce* complexity, but rather more critically: to help *manage* complexity.

Two important recent developments in related technologies will likely be central in this evolution: *Agentic AI* (AAI) platforms (WANG, L. et al. 2024) and Digital Twins, especially *Digital Twin Earths* (DTEs) (GRIEVES & VICKERS 2017).

These two emergent technologies promise to be mutually supporting and truly transformational for digital landscape architecture in the rest of the 21st century.

1.1 Agentic AI (AAI)

The 'chat'-style Generative AI many people are most familiar with now is a highly episodic interaction, initiated by a specific prompt, which elicits a response, and then often a refined or follow-on prompt spurred by some deficiency in the response, or some new avenue of inquiry, and so on. The human is not just in the loop but is leading the interaction at each step; the AI is mostly responding (or sometimes leading with canned "Do you want me to provide or explain anything in more detail?" questions). And there is effectively just one AI in the loop – typically, a LLM 'foundation model' trained on some very broad corpus of training samples. By contrast, a*gentic systems,* given high-level goals and other directives in initial prompts, can become proactive, goal-driven, and collaborative, using generative AI and other special-purpose software tools and APIs as components in a larger, self-directed, multi-step process. These agents may also be aware of change, and so can monitor changing contextual information, whether from sensors, or websites, or other sources. AAI systems aggregate several individual AI agents (AIAs) together, in a communicating network, often along with an 'orchestrator' or dispatcher agent (ANTHROPIC 2025).

Two additional developments promise to further control AI generative capabilities, and reduce hallucination, by adding an element of algorithmic/procedural workflow control via Model Context Protocol (MCP) (ANTHROPIC 2024); and factual lookup from one or more specialized / tailored databases via Retrieval-Augmented Generation (RAG), (LEWIS et al. 2020) which can provide the independent agents each with a suitable database of domain-specific information for their intended specialization. RAG-based agents can be much simpler than the full-bodied LLM behind generative AI; and their performance can be guided by rules, algorithms, and workflows managed by MCP, which can function like conventional computer code, running procedures, writing to databases, sending emails, and querying live APIs, rather than just probabilistic pattern matching.

1.2 Digital Twins (DTs)

Digital Twins are virtual replicas of real-world systems, capable of reflecting and predicting system behavior, such that ordinary and unusual conditions can be simulated computationally, monitored, and analyzed, and used to predict system behavior over time, under stress, etc. (WANG et al. 2024, HAZELEGER et al. 2024). First named and formalized in the context of modern aircraft design, construction, and operation, the idea has been spread to other domains and scales – from the human body, or individual organs therein, to the building-, region-, watershed-, and even planetary-scale – including the specially named Digital Twin Earths (DTEs). Perhaps the most advanced such efforts are Europe's Destination Earth (DestinE) (THE DIGITAL EUROPE PROGRAMME) and NASA's Earth System Digital Twins (ESDT) (ESA) initiatives; there are others, such as Google Earth Engine (GOOGLE 2024), from various sources (NASA), (NVIDIA) with various aims. DTEs can generate literally mountains of data, and so the interfaces for querying, analyzing, and storing their outputs is a rapidly evolving branch of technology. In more familiar traditional computer information terminology, AIAs contain *algorithms* or *procedures*, and the digital twins are databases – but in both cases, they are each hybrid, containing a combined mass of factual and procedural knowledge.

Specialized AIAs are especially well-suited to interface with special-purpose big data, such as with DTEs, giving rise to the following 'AAI+DTE' System proposal:

2 AAI+DTE Applied to Ervin's 'System for Geodesign'

A 'System for Geodesign', (ERVIN 2011) was originally presented at the DLA Conference as a proposal for a system of 15 distinct 'modules', which in combination would provide, in the author's view, the technological support – integrating CAD, GIS, database, version management, simulation, programming, and other functions – required for full-fledged geodesign as envisioned by (STEINITZ 2012). At the time, these modules were envisioned as tools such as computer users were accustomed to using; one at a time, each with limited and isolated functionality. This multi-module framing provides a useful test case to explore the design and eventual implementation of a suitable Agentic design system, and its AI Agents. Below, with some modifications to the number, order and the names as originally presented, are descriptions of 16 proposed Agents in the Agentic AAI+DTE approach:

(*Note that this is still mainly a thought-experiment and proposal; these prescriptions for agentic structure and behavior use 'could' and 'can', for now, rather than 'do' or even 'will'. An exploratory prototype instantiation of this outline proposal is described in the next section.)

1. Environment/Context-Base: The BaseMapAgent

With access to a robust DTE, the BasemapAgent could retrieve relevant layers for the location (s) and scale (s) required; and identify relevant contextual information. The BasemapAgent could automatically ingest, align, and integrate new data sources (imagery, sensor data, maps, etc.), even in real-time, working together with the Library Agent (#14).

2. Elements (Objects, Groups, Concepts): The ElementAgent

This agent maintains the working list of design elements, both concrete and abstract; and could propose site-appropriate landscape elements (landforms, structures, vegetation, paving, furniture, irrigation systems, etc.) based on context and goals, with reference to relevant libraries and catalogs. AI image recognition agents can classify and tag objects from aerial imagery, lidar, or text descriptions – enriching the elements database for design use.

3. Configuration: The ConfigAgent

This agent, using generative approaches (e. g. diffusion or GAN models) can create and modify alternative site configurations (paths, layouts, plantings, floorplans, etc.) responding to stated requirements or preferences. The ConfigAgent working with the ConstraintAgent could create, compare, and propose multiple feasible "what-if" scenarios. Note: This is still aspirational; such graphics capabilities and spatial reasoning are currently the least well-developed of AI technologies, and will require serious development work.

4. Constraints: The ConstraintAgent

The ConstraintAgent could receive high-level constraints directly from the human designer, as well as mine policy documents, local regulations, and site analysis to interpret and apply relevant constraints automatically. It could detect conflicts, and alert the human in the loop, or mediate between conflicting constraints (e. g., ecological vs. recreational goals), recommending optimized or satisficing resolution strategies.

5. Analyses: The AnalysisAgent (s)

This agent could conduct a wide range of geospatial and performance analyses, from standard GIS operations to more complex assessments like visual quality, habitat suitability, microclimate modeling, etc. It could proactively identify and run relevant analyses based on the design context and goals, highlighting potential issues for the human designer without needing to be explicitly prompted for each test.

6. Simulations: The SimAgent

The SimAgent could provide experimental dynamic analyses of proposed configurations and elements; using the resources of digital twins and procedures such as Computational Fluid Dynamics (CFD) and others to model dynamic processes over time, such as water flow, vegetation growth, energy performance, or pedestrian movement, etc. It could run multiple parallel simulations to compare the performance of different design alternatives under various future scenarios, using predictive models to forecast long-term ecological and social outcomes; as well as producing animations and other graphic artifacts.

7. Dashboards: The DashBoardAgent

This agent could generate and manage dynamic, real-time visualizations of the design's key performance indicators (KPIs). The DashBoardAgent could learn which KPIs are most important to the design team and create customized dashboards that clearly communicate progress towards project goals, working with the Abstraction and Diagram agents, and pulling data from both the Analysis and Simulation agents.

8. Version Manager: The VMAgent

The VMAgent, using version management technologies developed in the software industry, could track and manage the evolution of the design, capturing not just incremental saves but also significant branching points and design decisions. It could compare different versions, highlighting substantive changes in form or performance, and could suggest when prior abandoned solutions might have content applicable to a current design problem.

9. Time/Dynamics Manager: The TimeAgent

This agent would be responsible for modeling and integrating time-related aspects of the design. It could simulate and visualize phenomena across multiple temporal scales, from the daily cycle of light and shadow to seasonal vegetation changes and the long-term impacts of weathering or climate change. The TimeAgent would need to work closely together with the SimAgent.

10. Level of Abstraction (LOA) Manager: The AbstractionAgent

The AbstractionAgent would manage the various phases of the design project, and their associated representations, from the most abstract (initial programs, text representations, and schematic bubble diagrams, e. g.) to most concrete (3D models, working drawings, and detailed specifications). It would be aware of the stated design goals and intentions, and track the implications of those on specific, concrete proposed elements, layouts, and instantiations, and vice-versa. It can be used to help ensure that changes made at one level of detail are appropriately propagated to all other levels.

11. Diagram Manager: The DiagramAgent

Calling upon generative AI, the DiagramAgent could automatically translate complex data and spatial relationships into diagrams, charts, and conceptual graphics, for presentation to external decision-makers or on the design Dashboard. This agent could assist in communicating the design intent and performance to stakeholders, generating visual explanations for system dynamics or analysis results. Note: as mentioned above, as of this writing (2025), Gen AI is rather good at generating vivid and compelling illustrations, but is not very good at all at making clear diagrams, working drawings, or useful visual explanations – basic flowcharts are about as good as they get, and even then they are not yet always coherent; this remains an area for deep research and innovation.

12. Algorithmic Interface: The CodeAgent

The CodeAgent would provide a natural language interface for scripting and algorithmic design, allowing designers to generate custom functions without writing code manually. It could also respond to requests for code solutions from other agents and proactively search for and suggest novel algorithms or computational techniques from external libraries to solve unique or complex design challenges**.** In contrast to #11 above, this is a task that GenAI is already very good at.

13. Text/Media (Hyper-Annotations): The MediaAgent

This agent would link design elements to a rich network of related information, such as internal and world-wide web documents, including reports, images, and videos. It could automatically scan and tag documents for relevance and create hyper-annotations within the design model, providing contextual information built into the design documentation, so that the Dashboard manager could for example have multiple live links for deeper dives into relevant information.

14. Library: The LibAgent

The LibAgent would curate and manage extensive libraries and catalogs of design components, precedents, materials, and details, providing search and recommendation functions based on project context and constraints. Furthermore, it could provide materials and exemplars to generative models to create novel, context-appropriate design elements. The LibAgent would need to work closely with the ConfigAgent and ElementAgent, offering up elements and configurations from the Library, and storing newly developed or used ones for future use.

15. Collaboration tools: The CollabAgent

This agent would be responsible for facilitating teamwork among human designers, other stakeholders, and other AI agents, managing workflows, tracking tasks and calendars, and offering possible collaboration tools. It could mediate potential conflicts by summarizing differing viewpoints, highlighting areas of consensus, and proposing integration strategies for divergent ideas. This too is an area where AI is already quite functional.

16. Design Methods Coach: The DMCoachAgent

The DMCoachAgent would observe the overall design workflow and offers expert guidance on the process itself. It could recommend appropriate design methods, suggest timely analyses, and provide meta-level feedback on the team's progress. Depending on the state of the Dashboard, new agents may need to be acquired or spawned, and various functioning agents could be analyzed for their effectiveness, and replaced or re-written, by the DMCoachAgent and CodeAgent working together, according to some algorithm for evaluating ‘goodness’, ‘utility’, or ‘progress’, etc. Over time, the DMCoach could learn from experience, evaluation, and recording, what works or has worked well in what situation (s), building a library of ‘Learned Design Intelligence’.

3 Discussion: How Might These Technologies Work Together?

This proposal for a multipartite system is basically a reiteration of a familiar design office structure – a combination of experienced senior human designers and leaders, supported by varying numbers of junior staff and consultants. Different AI agents need not only different knowledgebases, priorities and precedents, but also different worldviews, procedures and algorithms. Although the above prescription identifies 16 variants, their exact number, makeup, training, and specializations will need to change and evolve with contexts, programs, client needs, and emergent priorities. For example, there is no 'ClimateMitigation' expert, or 'Sustainability' agent defined; this proposed system is at one abstraction level up from that. Climate mitigation criteria, plant selection, and other thematic or needs-based design criteria, will need to be individually addressed per design or program, by elements, configurations, analyses, simulations, collaborations, etc. However, one or more new generic agents may yet be required; for example, a dedicated 'Conductor' agent is likely to be essential in building an effective AIA design system (Fig. 1).

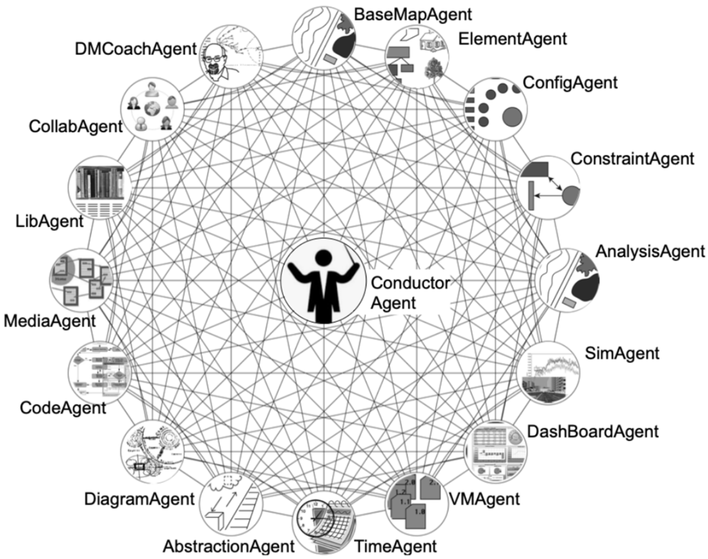

Fig. 1: 16 Distinct Generic Interconnected AI agents, all connected to a central Conductor agent

Design thinking and design processes have always been characterized by a tension, and some amount of cross-fertilization, between creative/imaginative/qualitative/intuitive on the one hand and'systematic/quantitative/evidence-based/rational on the other (the latter often supported by the engineering disciplines, the former by artistic practices); as well as between general/universal and site-specific/customized ideas and approaches. This can be seen already as a kind of *co-design*; the proposed AAI+DTE agentic system further crystallizes these combinations and competing/collaborating forces.

The multipartite structure of the proposed agentic system is designed to help manage complexity, by a combination of multiple agents and multiple specializations, just as in a big international, interdisciplinary design & construction firm; but soon accessible to a small 3-person firm, as well. This promise will depend very much on the development of specialized AI agents from within landscape architecture and allied disciplines, such as civil engineering, hydrology, conservation, forestry, transportation, etc.

Digital Twin Earths, which to date are mostly government-sponsored, are under active development. Connecting CAD, BIM and Agentic AI systems to their outputs from active sensors and archival databases is still a daunting challenge for designers. But this too will continue to evolve, and the access will inevitably become easier, just as the size and scope of data available will expand. Access to real-time data will increasingly become a part of responsive and responsible landscape design and will require the integration of specialized APIs and AI agents to achieve it.

Artificial (computerized) intelligence as a research and development effort is both broad and deep. At present, generative AI is still too prone to hallucinating (lying/fabricating) to be truly useful in any fact- or reality-based situation without a lot of human handholding. Because of the strictly statistical basis of its intelligence, it is usually plausible and sometimes even right, often enough to fool many into over-estimating its veracity and hence utility. There is real peril in this, and possible harm; and potential bias, political interference, and other ethical concerns plague even the largest training samples (NARAYANAN & KAPOOR 2024), (HANNA et al. 2025). But there's no denying that the combination of a huge accessible data base ('everything measurable on planetary sensor systems', or even 'everything on the World-Wide Web', e. g.), coupled with highly scalable pattern recognition and generative abilities, is a powerful force.

4 Proof-of-Concept Development: A Working AAI Prototype

In mid-November 2025, Google introduced AntiGravity (GOOGLE 2025), a new "agent-first" AI-powered Integrated Development Environment (IDE) for agentic software development. Using this system, the author has developed a modest but functional proof-of-concept prototype agentic system. This prototype represents a highly simplified instantiation of the proposed system architecture, implementing only three agents: a Conductor agent (orchestrating overall workflow), an Element agent (managing landscape components and materials), and a 'Sustainability' agent (not one of the core agents in the spec, but suggested by AntiGravity, and implemented as an experiment) – and all of those quite rudimentary. The system runs on a Mac Studio desktop computer using an open-source development stack (Python, PostgreSQL, and associated RAG and MCP libraries) with API connectivity to OpenAI's language models (ChatGPT-4o). Domain knowledge is provided through a small vector database populated with project examples scraped from a commercial landscape architecture office's public website, creating a minimal but authentic corpus of professional precedents for retrieval-augmented generation.

Despite its limited scope, this prototype successfully demonstrates the core principles of the agentic architecture proposed herein. The Conductor agent receives natural language design prompts, decomposes them into sub-tasks appropriate for the specialist agents, coordinates their activities, and synthesizes their outputs into conversational design suggestions. The Elements agent queries the vector databases to identify contextually appropriate plant materials, site furnishings, surface treatments, etc., based on project requirements and precedent similarity, while also generating suggestions without being specifically prompted. The Sustainability agent analyzes the proposal from the Elements agent and makes suggestions for alternatives. The interaction among these agents exhibits genuinely agentic behavior: the system maintains conversational context across multiple exchanges, requests clarifications, prompts for additional info, proposes alternatives based on retrieved precedents, and iteratively refines outputs in response to feedback.

As an example, consider a proposal for a public park with an emphasis on walking. The LibAgent agent might identify as a precedent the Imperial Garden at Tokyo, and its circulatory approach of a main path connecting entrances and exits, with subsidiary smaller paths branching off in loops. Passed to the AbstractionAgent, and the DiagramAgent, and with consultation from the BaseMapAgent, this could generate a simple circulation diagram; and from this the ConfigAgent could identify suitable entrance and exit locations. The ElementAgent could suggest plant materials, benches and lighting fixtures, which the ConfigAgent could locate.

Note: The creation, analysis and maintenance of graphical components – line diagrams, layout drawings, perspective renderings and the like – remain the most difficult and problematic for the current generation of AI tools, but there is no doubt that these limitations will begin to be overcome in the next few years, especially with impetus from disciplines and projects outside of the purview of landscape design. This – real effective ‘CAD’-integration – is the development most still-needed for the full flowering of the agentic design system imagined and proposed herein. The simple prototype developed in AntiGravity can handle the logical flow and the textual narrative, but not yet the drawings – especially the diagrams!

The prototype's technical implementation highlights both the current feasibility and the many remaining challenges of building domain-specific agentic systems. On one hand, contemporary tools make it surprisingly easy for a single developer to assemble a working multi-agent system in a matter of weeks. The combination of capable foundation models (LLMs accessed via API), vector databases for semantic retrieval, generative on-line systems like those from OpenAI and others, and orchestration frameworks like Antigravity dramatically lowers the barrier to entry for experimental agentic applications. On the other hand, scaling up this minimal three-agent system prototype will require substantially more time and expertise. The vector database, while demonstrating in principle the viability of RAG for domain-specific design knowledge, remains small and narrow; connecting to comprehensive Digital Twin Earth data sources, integrating simulation capabilities, and ensuring reliable performance across diverse project types will require substantially greater effort, infrastructure, and investment.

5 Conclusion

This paper began as an effort to reconceptualize the 2011 'System for Geodesign' in light of transformative developments in artificial intelligence, geospatial technology, and environmental data infrastructure. The resulting proposal – a multi-agent Agentic AI system integrated with Digital Twin Earth resources – represents a speculative vision, arguing that such an architecture could address the escalating complexity, data intensity, and collaborative demands of contemporary landscape architecture practice.

The development of a working prototype, however modest, is a step from the purely hypothetical toward the increasingly feasible. The ability to build this prototype on Google's Antigravity platform also reveals something important about the trajectory of AI development. Antigravity is itself an Agentic AI system – a generalized platform for orchestrating multiple specialized agents toward complex project goals – that happens not to be domain-specific to landscape architecture. Creating a landscape-focused agentic system within Antigravity demonstrates a kind of fractal recursion: an agentic architecture building an agentic architecture, with the same organizational principles applying at both the development-tool level and the application level. This pattern is not unique to Antigravity; similar multi-agent orchestration frameworks are emerging across the AI development ecosystem, suggesting that the agentic paradigm is becoming a standard architectural pattern for complex AI-enabled work (and, as a side-effect, with fewer 'humans in the loop' of coding and software development!)

The challenges ahead remain substantial. Scaling from three agents to sixteen, from a small precedent corpus to comprehensive integration with Digital Twin Earth resources, from schematic suggestions to full design documentation and effective CAD-integration, and from single-user prototypes to collaborative professional tools, will require significant technical development, disciplinary knowledge formalization, and careful attention to the practical and ethical dimensions of AI-augmented design practice. One of the more intriguing, and challenging aspects of the proposed system is the ability for it to learn over time, not just about elements and landscapes, but also about patterns of design strategies, building up a database over time of 'Learned Design Intelligence' from its use.

The landscape architecture profession has absorbed and responded to technological shifts for decades; the prototype described here and the platforms on which it was built are just the latest stage in the evolution of 'computer aided design'; and they are still fundamentally lacking. In spite of the much-publicized generative facilities with text and graphics, most AI foundation models are still quite primitive with respect to graphical representation and spatial reasoning – essential for landscape architects. Another principal missing ingredient in most current AI systems is algorithmic and rule-based logical reasoning and calculating components, that surely will necessarily underlie the ultimate goal of Artificial General Intelligence (AGI) (GOERTZEL & PENNACHIN 2007). This is the premise of the latest developments in *neuro-symbolic AI* (D'AVILA GARCEZ & LAMB 2023), (of which RAG is an early example) in which the neural-net pattern-recognition and statistical generative synthesis tools are combined with more algorithmic and rule- and logic-based symbolic approaches, including appropriate constraints and 'guard-rails', and so less prone to fabrication, and so more defensible and reliable. Current agentic platforms, development IDEs like AntiGravity, and RAG and MCP implementations are all still in their infancy. The next developments in this synthetic environment will provide even better support for next-gen computer-aided design, as design needs both rules and rule-breaking, both dreaming and solid structural foundations;

and effective performative designs will require access to real-world dynamic sensor data, and may be accomplished by robotic ‘actuators’ in the field.

Perhaps not all landscape designs or design processes need to be deterministic, just as not all landscape architects need to be digital landscape architects – but many of them, for infrastructure and other large and public works, benefit from being so. A well-formed neuro-symbolic AIA agentic system, with RAG/ MCP augmentation, connected to a database repository of reliable domain-specific information, and a robust DTE – along with one or more humans in the loop (Fig. 2), to provide ‘intentionality, contextual awareness, and critical reflection’, among other things – is a likely candidate for delivering such tools to the next generation of digital landscape architects and planners.

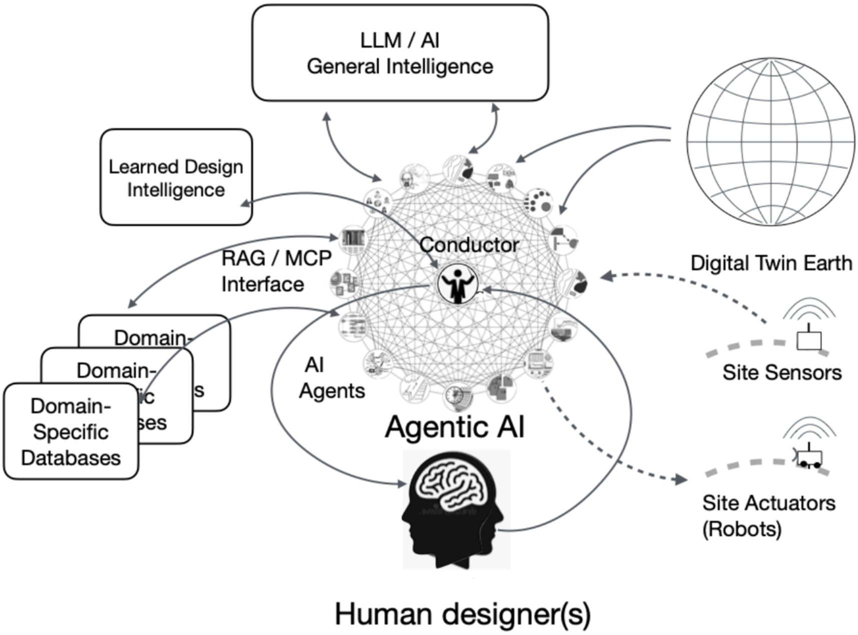

Fig. 2: Conceptual Architecture of the AAI+DTE System. This diagram illustrates the functional ecosystem of the proposed Agentic AI design system. At the center is the Agentic AI cluster, featuring the network of specialized Agents (e. g., Layout, Elements, Simulation) orchestrated by a central Conductor agent. The Human designer (s) ("human-in-the-loop") interface primarily with the Conductor to provide intent, goals, and critical review. The agents’ cognitive capabilities are powered by LLM / AI General Intelligence (foundation models), supplemented by Learned Design Intelligence (specific design heuristics and rules). To ensure accuracy and reduce hallucination, the system utilizes a RAG Interface (Retrieval-Augmented Generation) and Model Context Protocol (MCP) library to provide algorithmic flow-control and to fetch verified information from Domain-Specific Databases (containing precedents, codes, and material libraries). Finally, the design is grounded in real-world context through connections to a Digital Twin Earth for environmental data; and Site Sensors / Actuators, not addressed in this paper, can provide on-site monitoring and machine-controlled (‘robotic’) physical interventions.

References

ABIOYE, S. O., OYEDELE, L. O., AKANBI, L. A., AJAYI, A. O., DELGADO, J. M. D., BILAL, M. & AHMED, A. (2021), Artificial intelligence in the construction industry: A review of present status, opportunities, and future challenges. Journal of Building Engineering, 44, 103299.

ANTHROPIC (2024), Building effective agents. https://www.anthropic.com/research/building-effective-agents (24.10.2025).

ANTHROPIC (2025), Introducing the Model Context Protocol. https://www.anthropic.com/news/model-context-protocol (10.01.2026).

BOWEN, J. A. & WATSON, E. (2024), Teaching with AI: A Practical Guide to a New Era of Human Learning. Johns Hopkins University Press. https://doi.org/10.56021/9781421449227.

BRAIDEN, H., CHAMBERLAIN, B., GEORGE, B. H. & FERNBERG, P. (2025), AI in Practice: Professional Survey Findings from Landscape Architects in North America. JoDLA, 10 (2025), 626-633.

CANTRELL, B., ZHANG, Z. & LIU, X. (2021), Artificial intelligence and machine learning in landscape architecture. In: As, I. & Basu, P. (Eds.), The Routledge Companion to Artificial Intelligence in Architecture. Routledge.

D'AVILA GARCEZ, A. & LAMB, L. C. (2023), Neurosymbolic AI: the 3rd wave. Artificial Intelligence Review, 56 (11), 12387-12406.

ERVIN, S. (2011), A System for GeoDesign. In: Proceedings of Digital Landscape Architecture. Anhalt University of Applied Science.

EUROPEAN SPACE AGENCY (2025), ESA's Digital Twin Earth programme: building a virtual model for a changing planet. https://www.esa.int/Applications/Observing_the_Earth/ESA_s_Digital_Twin_Earth_programme_building_a_virtual_model_for_a_changing_planet (24.10.2025).

GOERTZEL, B. & PENNACHIN, C. (Eds.) (2007), Artificial General Intelligence. Springer.

GOOGLE (2024), Google Earth Engine. https://earthengine.google.com (29.12.2025).

GOOGLE (2025), Antigravity. https://antigravity.google (10.01.2026).

GRIEVES, M. & VICKERS, J. (2017), Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems. In: Transdisciplinary Perspectives on Complex Systems. Springer.

HANNA, M. G., PANTANOWITZ, L., JACKSON, B., PALMER, O., VISWESWARAN, S., PANTANOWITZ, J., DEEBAJAH, M. & RASHIDI, H. H. (2025), Ethical and Bias Considerations in Machine Learning. Modern Pathology, 38 (3).

HAZELEGER, W., AERTS, J. P. M., BAUER, P., BIERKENS, M. F. P., CAMPS-VALLS, G., DEKKER, M. M., DOBLAS-REYES, F. J., EYRING, V., FINKENAUER, C., GRUNDNER, A., HACHINGER, S., HALL, D. M., HARTMANN, T., IGLESIAS-SUAREZ, F., JANSSENS, M., JONES, E. R., KÖLLING, T., LEES, M., LHERMITTE, S., VAN NIEUWPOORT, R. V., PAHKER, A.-K., PELLICER-VALERO, O. J., PIJPERS, F. P., SIIBAK, A., SPITZER, J., STEVENS, B., VASCONCELOS, V. V. & VOSSEPOEL, F. C. (2024), Digital twins of the Earth with and for humans. Communications Earth & Environment, 5, 463. https://doi.org/10.1038/s43247-024-01626-x.

LEWIS, P., PEREZ, E., PIKTUS, A., PETRONI, F., KARPUKHIN, V., GOYAL, N., KÜTTLER, H., LEWIS, M., YIH, W.-T., ROCKTÄSCHEL, T., RIEDEL, S. & KIELA, D. (2020), Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. Advances in Neural Information Processing Systems (NeurIPS), 33, 9459-9474.

MINSKY, M. (1986), The Society of Mind. Simon & Schuster.
NARAYANAN, A. & KAPOOR, S. (2024), AI snake oil: what artificial intelligence can do, what it can't, and how to tell the difference. Princeton University Press.
NASA (2025), Earth System Digital Twins (ESDT). https://esto.nasa.gov/earth-system-digital-twin/ (24.10.2025).
NVIDIA (2025a), Earth-2. https://www.nvidia.com/en-us/high-performance-computing/earth-2/ (24.10.2025).
NVIDIA (2025b), Retrieval-Augmented Generation. https://www.nvidia.com/en-us/glossary/retrieval-augmented-generation/ (29.12.2025).
SCHÖN, D. (1983), The Reflective Practitioner: How Professionals Think in Action. Temple Smith.
SIMON, H. (1996), The Sciences of the Artificial (3rd Ed.). MIT Press.
STEINITZ, C. (2012), A Framework for Geodesign: Changing Geography by Design. Esri Press.
THE DIGITAL EUROPE PROGRAMME (DestinE) (2025), https://destination-earth.eu (24.10.2025).
TURING, A. (1948), Intelligent Machinery. In: Copeland, B. J. (Ed.), The Essential Turing. Oxford University Press. https://doi.org/10.1093/oso/9780198250791.003.0016 (24.10.2025).
WANG, L., MA, C., FENG, X., ZHANG, Z., YANG, H., ZHANG, J., CHEN, Z., TANG, J., CHEN, X., LIN, Y., ZHAO, W. X., WEI, Z. & WEN, J.-R. (2024), A Survey on Large Language Model based Autonomous Agents. Frontiers of Computer Science. https://arxiv.org/abs/2308.11432.
WANG, Y., WANG, L., HAN, W. & YAN, J. (2024), Digital twin of Earth: A novel information framework for managing a sustainable Earth. The Innovation Geoscience, 2 (4), 100092. https://doi.org/10.59717/j.xinn-geo.2024.100092.
WIKIPEDIA (2025), Human-in-the-loop. https://en.wikipedia.org/wiki/Human-in-the-loop (24.10.2025).
XING, Y., GAN, W. & CHEN, Q. (2025), Artificial intelligence in landscape architecture: A survey. International Journal of Machine Learning and Cybernetics, 16, 4685-4710. https://doi.org/10.1007/s13042-025-02536-w.