LAF | https://doi.org/10.15302/J-LAF-2026-0032
Volume 14 / Issue 4 / 260032 / 2026
Received Date | 2026-04-13 Accepted Date | 2026-05-27
© Higher Education Press 2026

Interpolative and Extrapolative Work —The Place of Human Intelligence in the Age of Artificial Intelligence

Liyan XU*

College of Architecture and Landscape, Peking University, Beijing 100080, China

*CORRESPONDING AUTHOR
Address: No. 3 North Street, Yandongyuan, Haidian District, Beijing 100080, China
Email: xuliyan@pku.edu.cn

The life of the mind, whether in scientific research or design practice, falls broadly into two kinds of work. One operates within the boundaries of what is already known—refining established patterns, filling in details, and completing the picture. We may call this interpolative work. The other ventures to the edges of human knowledge, posing new questions and forging connections that did not previously exist. This we may call extrapolative work[1].

The distinction is hardly new. But the rapid advance of artificial intelligence (AI) has transformed it from an epistemological nicety into an urgent practical concern. In recent years, generative AI has made striking progress across text, images, code, and other modalities. Frontier foundation models display not only remarkable breadth of knowledge but also considerable depth of reasoning, performing at levels that rival experienced researchers on many tasks[2]. In the design disciplines, vision-language models (VLMs) have demonstrated impressive powers in visual composition, achieving high standards in creative generation and visual representation. Combined with auxiliary tools such as autonomous agents, contemporary AI systems now possess formidable working capabilities—enough to provoke widespread anxiety. Yet before we give in to that anxiety, perhaps we should try to understand, with some composure, what exactly is AI? What can it do, and what can it not do? And what work remains ours if we must coexist with AI?

The author would like to advance the following proposition: current AI is, at its core, a powerful interpolation engine. This implies that any work that is highly procedural and can be carried out within existing knowledge frameworks, however sophisticated it may appear at present, will eventually be subject to automation by AI. Extrapolative work, by contrast, constitutes the last sustainable stronghold of human intelligence. This is at once a claim about the limitations of AI and a recognition of our own responsibility.

To understand why AI functions as an interpolation engine, it is worth returning to first principles. The attention mechanism[3] that underpins contemporary large language models (LLMs) is mathematically isomorphic to the Boltzmann distribution in statistical mechanics[4]. The two share an identical exponential normalization structure, to the extent that even the hyperparameter governing the shape of the distribution, namely the physicist's notion of "temperature," has been directly inherited by AI systems. While the correspondence does not imply that the two are equivalent in any physical sense, neither is it a purely formal coincidence: it means that at each step of computation, the attention mechanism seeks the statistically optimal weighting given the context—a local, training-data-based state of probabilistic equilibrium. Generalizing this local mechanism to the generative process as a whole, the workings of current AI can be understood, to a reasonable approximation, as a probabilistic completion system: given an input (i.e., a prompt), it searches the probability space spanned by its training data for the statistically most likely output.

This observation carries at least four successive implications for the question of AI's capability boundaries. At the most immediate, data level, AI performs brilliantly within the domain covered by its training data; by definition, this is interpolation. Yet when a task extends significantly beyond the stable regions supported by prior

training experience, the model tends to resort to completions that are statistically plausible yet factually groundless. Because of the compulsory nature of Softmax normalization, such an output is produced regardless, giving rise to what is commonly termed "hallucination"[5]. In other words, current AI excels at identifying the most probable extensions of existing statistical structures, rather than creating genuinely new problem spaces. Two examples from the discipline of Landscape Architecture may help make this point concrete. The first concerns architectural and landscape renderings. The visual grammar of renderings, including perspective, materiality, light, and planting composition, has already been extensively covered by vast quantities of training data, enabling AI to generate high-quality images of any style or viewpoint with remarkable efficiency. The second concerns seemingly more "advanced" tasks, such as GIS-based site suitability analysis, carefully specified statistical models, and even certain forms of causal inference designs. Although these tasks still require considerable professional training at present, they remain, at bottom, the processing of data within established methodological frameworks—interpolation in essence, and thus in principle within the reach of AI. By contrast, design work grounded in genuine authorial intent, including distinctive problem consciousness, value positions, and conceptual orientations, still depends primarily on human judgement and experience—that is, on extrapolation.

Will the accumulation of data, cases, and experience, then, eventually diminish the value of human extrapolation? Not necessarily—for there exists a class of limitations that lies not in the quantity of data but in the very mechanism of the model itself. In design work, one readily notices that existing AI systems struggle with concepts of spatial configuration such as "rhythm" and "balance," and the root of such limitation lies in the basic architecture of the algorithm. Concepts of spatial configuration are, in essence, holistic (*Gestalt*) or topologically high-order: they depend on the complex global ordered relational structure among all objects in a scene, rather than on any statistical superposition of local features. The prevailing attention mechanism is built on pairwise similarity computation. While it does not in principle forbid the capture of such higher-order relations, the cost of learning them becomes prohibitively high in practice under realistic data and computational budgets. Here, because the problem lies in the structure of representation rather than in data coverage, simply increasing the volume of data cannot fundamentally resolve it. It is in this way that the traditional craft of design retains its irreplaceability. This is the second sense of extrapolation.

Going further, even when we confine our attention to the algorithmic mechanism itself, the physical picture implicit in the current architecture prescribes its inherent boundaries. Contemporary AI architecture corresponds roughly to the equilibrium statistical mechanics of the mid-to-late 19th century; the treatment of ordered structures far from equilibrium, and of problems entailing the arrow of time and endogenous causality, lies in principle beyond its scope. Equilibrium statistical mechanics describes the most probable state of a system under given constraints, and one of its defining features is history-independence: the system's present state is fully determined by its energy landscape, irrespective of how it arrived at that state. Endogenous causality, however, means that a system's past states exert an irreversible influence on its future—which is precisely the physical meaning of the arrow of time, and the very thing that distinguishes a dissipative system from one at equilibrium[6]. To break out of the existing architecture—by analogy, to move from the Boltzmann–Gibbs equilibrium toward a non-equilibrium dissipative system in Prigogine's sense—may require something of a paradigm shift. Yann LeCun's proposed energy-based architecture for autonomous intelligent agents[7] may be regarded as one of the most ambitious attempts to date; yet, as he himself acknowledged, the path remains at a highly exploratory frontier and cannot promise a breakthrough within any foreseeable horizon.

Finally—and most fundamentally—even should the difficulties above one day be overcome by technological breakthrough, AI's capacity for extrapolation would still face an epistemological constraint that is ineradicable and independent of any technical assumption. The early work of Ludwig Wittgenstein famously declared that the limits of language are the limits of the world[8]. If this proposition holds, then AI, confined as it is to the symbolic world of text, necessarily faces an insurmountable cognitive ceiling. In *Philosophical Investigations*[9], the later Wittgenstein pressed further into the source of this boundary: the meaning of language lies not in its logical mapping onto the world, but in the "form of life" (*Lebensform*) within which it is embedded. Meaning is use, and use is sustained by an entire fabric of embodied, social practice. If one accepts this insight, its delimitation of AI's capability becomes all the more decisive—especially for the discipline of design, for it implies that the experience of everyday life, design included, is ultimately irreducible to any reductionist account. A language model can learn every correct usage of terms such as "sense of scale," "genius loci," and "spatial sequence," and may even deploy these terms with apparent precision in the right context. Yet AI has no body. It

has never walked along a mountain trail, struggled through dense forest, and emerged onto an open ridge. Design is fundamentally a practice grounded in bodily experience and the experience of place: the sensation of earth and gravel underfoot, the angle of sunlight filtering through a canopy, the coolness of wind funneling up from a valley, the shifting rhythm of exertion during a long ascent, and the sense of release upon arrival. These experiences correspond precisely to what Wittgenstein described as a "form of life"—the largely inexpressible ground that sustains the meaning of design language. In any reasonable sense, this also belongs to the realm of extrapolation, lying well beyond the reach of mere "computation."

We should not, of course, claim that AI will never achieve embodied understanding. Under the current paradigm, however, this limitation appears structural rather than accidental. Infants offer a telling illustration: although incapable of speech, they plainly possess intelligence—an intelligence grounded in the body, perception, lived experience, and interaction, and one that precedes language. Designers, perhaps, will recognize this observation readily. The beginner's mind, or what Chinese thought describes as the "heart of the newborn," remains among the most valuable qualities in design and design research alike.

If the foregoing argument is accepted, then rather than fretting over replacement by AI, it may be more constructive to ask what role human beings should assume within a new division of intellectual labor. A pragmatic principle suggests itself: delegate interpolative tasks to AI while concentrating human effort on the extrapolative frontier.

For researchers, AI will grow increasingly capable of undertaking tasks that are methodologically executive, informationally integrative, or rhetorically routine, such as literature synthesis, data processing, model fitting, and the drafting of conventional sections of academic writing. These tasks are undoubtedly important and often constitute the foundation of research. Yet insofar as they operate within existing problem framings and established methodological repertoires, their cognitive character remains closer to interpolation than to extrapolation. Researchers' irreducible responsibility thus converges on a single task: identifying genuinely valuable research questions and articulating a principled approach to addressing them. The formulation of meaningful questions is inherently extrapolative because it often requires establishing previously nonexistent connections across knowledge domains—connections that lie outside the probability space of any existing training corpus.

For designers, AI offers a virtually infinite space of creative possibilities, including form generation, scheme variation, and scene rendering. Yet designers' role is not to select blindly from this space, but to define its direction with intention: to determine what is worth doing, why it is worth doing, and for whom. At the same time, those aspects of design work that require bodily presence, such as sensing a site, empathizing with users, communicating with clients, and constructing narratives for the public, also belong to the broader domain of extrapolation. For this reason, they will remain, and are likely to remain for a long time, fundamentally human responsibilities.

We may continue with the earlier example to further illustrate the boundary of such human–AI collaboration. Consider the task of planning and designing an outdoor trail system. AI can mine vast repositories of hiking GNSS traces, identifying statistical patterns and answering, with reasonable rigor, questions such as which segments are the most popular, what kinds of users prefer what kinds of routes, and what the optimal path combination might be under a given fitness budget. All these tasks are completion within the probability space of existing behavioral data and therefore belong paradigmatically to interpolation. But judging the experiential quality of a trail, such as the shifting rhythm of ascent, the dramatic opening of a view around a bend, the alternating cadence of shade and exposure, and the sense of fulfillment experienced when looking back over the route from the summit, requires perceptual capacities grounded in bodily presence. In this respect, such judgements remain within designers' irreplaceable domain. Excellent work should therefore synthesize both dimensions: AI providing an efficient data infrastructure, and designers endowing the routes with an experiential meaning structure.

Placed within a broader historical perspective, we should recognize that the standard for what counts as intellectually creative research has always been in flux[10]. Descriptive statistical reports, for example, were regarded as serious scholarly contributions in the 19th century; today, few would view them as such. In landscape architecture, rendering was once considered a meaningful and valuable professional task; yet it, too, is rapidly being replaced by automation. The rising standards might simply be a natural consequence of broader advances in knowledge and capability. More fundamentally, however, as the knowledge frontier advances and tools become increasingly capable, forms of work that once occupied the frontier are gradually absorbed into the interpolable—first by human effort, then by machine. AI has not altered the direction of this process; it has merely accelerated it enormously.

From this, we might venture a more general criterion: if a task

can be described as the application of a known methodological framework to new data or a new context, then it is interpolative in nature, regardless of how technically intensive it may appear at present. This should not be understood dismissively. As noted earlier, interpolative work has genuine practical value. Yet its academic lifespan is ultimately finite. By contrast, extrapolative work begins precisely where existing knowledge cannot automatically generate answers—where one must assume responsibility for setting direction, making value judgements, and conferring meaning through experience.

What, then, constitutes genuinely extrapolative work in the context of landscape architecture and its research? Is it the invention of entirely new principles of spatial organization? The discovery of previously unknown mechanisms in human–environment interaction? Or the articulation of new value judgements about what constitutes a good landscape? These questions have no ready answers. But that is precisely the point. To formulate a genuinely important question is itself the purest form of extrapolation.

For current professionals and students alike, transitioning to extrapolative work will naturally not be an easy task. Beyond the heightened intellectual demands, the impact of the new division of labor on the field's existing structure may prove an even more pressing tension. Renderings again offer a revealing example. We must recognize that it is not merely a professional task but also a formative training path through which young designers cultivate their spatial intuition. When AI "resolves" this task, what disappears is not only a category of employment, but potentially an entire apprenticeship route toward design sensibility. The challenge is real and profound; perhaps this itself constitutes another "extrapolative" problem left to us by the present era.

In the long run, AI's capabilities will undoubtedly continue to expand, perhaps far beyond what we can imagine today. But even so—or rather, precisely because of this—the pursuit of extrapolative research and design at the known frontier is arguably the one intellectual commitment that is uniquely and most worthily human. As the scientific community wrote 80 years ago: the frontier is endless[11]. In the age of AI, we should regard this not merely as a piece of rhetoric. More than ever, it is an invitation.

Competing interests | The author declares that he has no competing interests.

REFERENCES

[1] Chollet, F. (2019). On the measure of intelligence (arXiv:1911.01547). *arXiv*.

[2] Bubeck, S., Chandrasekaran, V., Eldan, R., Gehrke, J., Horvitz, E., Kamar, E., Lee, P., Lee, Y. T., Li, Y., Lundberg, S., Nori, H., Palangi, H., Ribeiro, M. T., & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4 (arXiv:2303.12712). *arXiv*.

[3] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *arXiv*.

[4] Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical Review*, *106*(4), 620–630.

[5] Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, *35*, 27730–27744.

[6] Prigogine, I., & Stengers, I. (1984). *Order Out of Chaos: Man's New Dialogue With Nature*. Bantam Books.

[7] LeCun, Y. (2022). *A path towards autonomous machine intelligence* (Version 0.9.2).

[8] Wittgenstein, L. (1921). *Tractatus Logico-Philosophicu*s. Kegan Paul.

[9] Wittgenstein, L. (1953). *Philosophical Investigations*. Blackwell.

[10] Kuhn, T. S. (1997). *The Structure of Scientific Revolutions*. University of Chicago Press.

[11] Bush, V. (1945). *Science: The Endless Frontier*. U.S. Government Printing Office.

LAF | https://doi.org/10.15302/J-LAF-2026-0032
第14卷 / 第4期 / 260032 / 2026
收稿日期 | 2026-04-13　录用日期 | 2026-05-27
中图分类号 | TU984　文献标识码 | B

内插性工作与外推性工作——AI时代的人类智识定位

许立言*

北京大学建筑与景观设计学院，北京 100080

*通信作者
地址：北京市海淀区燕东园北小街3号
邮编：100080
邮箱：xuliyan@pku.edu.cn

人类的智识活动，无论在科学研究还是设计创作中，大体可分为两类：一类耕耘于既有知识的边界之内，细化已有模式，补全各种细节，不妨称之为“内插性”（interpolative）工作；另一类则游走于人类知识的边缘，提出新的问题，建立此前不存在的联系，不妨称之为“外推性”（extrapolative）工作[1]。

内插与外推的二分法并不新鲜，但人工智能（AI）的迅猛发展使其从一个认识论问题转变为一个紧迫的实践问题。近年来，生成式AI在文本、图像、代码等多个模态上取得了显著进展。前沿基础模型不仅展现出广泛的知识广度，也表现出相当的推理深度，其水平在许多任务上已不逊于资深研究者[2]。在设计领域，视觉语言模型（VLM）显示了令人惊艳的图像构成能力，在创意生成与效果表达环节达到了相当高的水平。结合自主智能体等辅助工具，当代AI系统已经形成较强的实际工作能力，并由此引发了广泛的焦虑。不过，焦虑之余，我们似乎应尝试冷静地理解：AI究竟是什么？它能做什么、不能做什么？以及，如果人类必须与AI共存，留给我们的工作还有什么？

笔者谨提出以下观点：当前AI本质上是一个强大的内插引擎。这意味着，凡属高度程式化、能够在既有知识框架内展开的内插性工作，无论其当前表现得多么精致，最终都可能被AI的自动化替代；而外推性工作则构成人类最后且可持续的智识根据地。这既是关于AI能力边界的判断，也是关于人类自身责任的认知。

为理解AI为何属于内插引擎，我们不妨回到基本的数学原理。当前主流大语言模型的基石是注意力机制[3]，它在数学形式上与统计力学中的玻尔兹曼分布是同构的[4]。二者共享完全相同的指数归一化结构，乃至用于控制分布形态的超参数（即物理学中的“温度”概念）也被AI系统原封不动地继承了。尽管这种对应关系并不意味着二者物理意义的等同，但它也绝非纯粹形式上的巧合：它意味着注意力机制在每一步计算中寻找的都是给定上下文中统计意义上的最优加权——一种局部的、基于训练数据的概率平衡状态。将此局部机制推广到整个生成过程，当前AI的工作方式可被近似理解为一种概率补全系统：给定输入（即提示词），模型将在训练数据所张成的概率空间中，寻找统计意义上最可能的输出。

以上观察引出了关于当前AI能力边界问题的至少4层意涵。在最直接的数据层面上，AI在训练数据所覆盖的范围内表现卓越，这在定义上正属于“内插”。然而，当任务明显超出其训练经验所支撑的稳定区域时，模型往往会诉诸统计上看似合理、实则缺乏依据的补全。由于Softmax归一化的强制性特征，此类结果仍将被输出，由此产生所谓“幻觉”[5]。换言之，当前AI所擅长的是在既有统计结构中寻找最可能的延展，而不是实质性地发明新的问题空间。在景观设计学科语境下，以下两个例子或可直观说明这一点。其一是建筑与景观效果图制作。效果图中的视觉语法（如透视、材质、光影与植物配置等）已被海量训练数据充分覆盖，因此AI能够快速批量生成任意风格与视角的高质量渲染图。其二则涉及看似更为“高级”的任务，例如应用了GIS的场地适宜性分析、复杂统计模型构建，乃至某些因果推断设计。这些工作目前仍需较高程度的专业训练，但终究属于在既有方法框架内对数据进行处理——本质上仍是内插，原则上仍处于AI可及的范围之内。相比之下，在问题意识、价值立场与概念方向上真正体现了作者性的设计工作，目前仍主要依赖人的判断与经验，即属于“外推”。

那么，数据、案例和经验的积累是否终将削弱人类外推的价值？未必。因为有一类限制并不在于数据的多寡，而在于模型的机制本身。在设计工作中，不难注意到现有AI系统难以处理诸如“韵律”或“均衡”

之类的空间构成概念，而这一限制的根源在于算法的基本架构。空间构成概念在本质上是整体完形（*Gestalt*）或高阶拓扑关系，依赖于场景中所有对象之间复杂的全域有序关系结构，而非任何局部特征的统计叠加。现有的注意力机制以成对元素的相似度计算为基础，原则上虽不禁止捕捉此类高阶关系，但在现实的数据规模与计算预算下，其学习成本往往高到难以实现。这里的问题出在表示结构而非数据覆盖范围，因此单纯增加数据量并不能根本性地解决此问题。传统的设计艺匠由此得以保留其不可替代性。这是外推的第二层意涵。

更进一步，即使着眼于算法机制本身，当前架构所隐含的物理图景也规定了其固有的边界。现有AI架构大致对应于19世纪中后期形成的平衡态统计力学框架，而对于远离平衡态的有序结构，以及涉及“时间之矢”和“内生因果性”等问题的处理，原则上超出了其范畴。平衡态统计力学描述的是系统在给定约束下的最可能状态，其核心特征之一是历史无关性：系统的当前状态由其能量格局完全决定，而与它如何到达这一状态无关。然而，内生因果意味着系统的过去状态对未来有不可逆的影响——这正是时间之矢的物理含义，也是耗散系统区别于平衡态系统的根本所在[6]。要突破现有架构，类比而言，即从玻尔兹曼—吉布斯平衡态走向普里戈金意义上的非平衡态耗散系统，可能需要某种范式转换。杨立昆提出的基于能量函数的自主智能体架构[7]，可被视为当前最具雄心的尝试之一；但正如他本人所承认的，该路径仍处于高度探索性的前沿阶段，尚不能承诺在可预见的未来取得突破。

最后，也是最根本的，纵使以上困难在未来某一天被技术突破所克服，AI的外推能力仍存在一个无从消解、不依赖任何技术假设的认识论约束。路德维希·维特根斯坦在早期曾提出一个著名的、当下也被反复引用的命题：语言的边界即为世界的边界[8]。如果该命题成立，那么受限于文本符号世界的AI，自然存在某种不可逾越的认知上限。晚期维特根斯坦在《哲学研究》中进一步追问了此边界的根源[9]：语言的意义并不在于对世界的逻辑映射，而在于它所嵌入的“生活形式”（*Lebensform*）。意义即用法，而“用法”则依赖于完整的身体性与社会性生活实践。如果接受这一洞见，那么它对于AI能力边界的限定将更具决定性——对设计学科而言尤其如此，因为这意味着包含设计在内的日常生活体验，从根本上无法以还原论的方式化约。语言模型或许能够学会“尺度感”“场所精神”“空间序列”等设计概念的全部正确用法，甚至可以在恰当的语境中貌似精确地使用这些术语；但AI并没有身体，它不曾真正走过一条山路，也不曾在密林跋涉后抵达豁然开朗的山脊。设计本质上是一种深深嵌入身体经验与场所经验的实践：脚踩泥土与碎石的触感、阳光穿透树冠的角度、山谷冷风扑面而来的凉意、长时间攀升后的体力节奏变化，以及抵达时的释放感。这些经验，恰恰构成了维特根斯坦所谓“生活形式”的一部分——即支撑设计语言意义的、难以完全言说的基底。从合理意义上讲，这也是外推，是单纯的“计算”所不能触及的领域。

当然，我们不应断言AI“永远”无法获得身体性的理解。但综上所述，这是一个结构性——而非偶然性——的局限。婴儿本身便提供了一个有力例证：他们尚不会说话，却显然已经具备某种智能——一种建立在身体、感知、体验和交互之上的智能，并且先于语言而存在。设计师对此或许尤能产生共鸣。所谓“赤子之心”，正是设计乃至设计研究中最宝贵的内核。

如上述判断成立，那么与其焦虑于被AI取代，不如认真思考一个更具建设性的问题：在新的智识分工之下，人类应当承担什么角色？一个务实的原则或许是：将内插性工作交给AI，而将人类的精力集中于外推性的前沿。

对研究者而言，AI将越来越胜任大量方法执行性、资料整合性和表达规范性工作，如文献梳理、数据处理、模型拟合，以及学术写作中的一般性段落撰写等。这些工作当然重要，甚至往往构成研究的基础；但只要它们主要在既有问题意识与方法框架内展开，其认知性质便仍更接近内插而非外推。因此，研究者不可替代的任务最终收束为一点：提出真正有价值的研究问题，并为其提供原则性解决思路。问题意识之所以具有外推性，正是因为真正重要的问题往往要求在不同知识域之间建立此前不存在的联系，而这种联系并不位于任何训练数据的概率空间之内。

对设计师而言，AI提供了一个近乎无限的创意可能性空间，包括形式生成、方案变体、场景渲染等。然而，设计师的职责并不是在这一空间中盲目挑选，而是自觉地规定其方向：决定什么值得做、为何值得做，以及为谁而做。同时，设计工作中那些需要身体在场的环节，包括对场地的感知、对使用者的共情、与业主的沟通，以及对公众的叙事等，同样属于广义的外推。因此，这些工作不仅仍将属于人类，而且在相当长时期内都难以被替代。

我们或可继续沿用前述案例，进一步说明此类人机协作的边界。以户外游线规划设计为例，AI能够从海量全球导航卫星系统（GNSS）徒步轨迹数据中识别统计模式，并相当可靠地回答诸如哪些路段最受欢迎、不同类型用户偏好何种路线，以及在既定体力预算下何种路径组合最优等问题。以上种种都是在既有行为数据的概率空间中做补全，因而属于典型的内插性任务。然而，对于一条游线体验品质的判断——例如攀升节奏的层次变化、转弯处视野豁然开朗的戏剧性、树荫与暴晒交替形成的韵律，以及抵达终点后回望来路时的满足感——则需要身体在场的感知能力，因此仍属于设计师不可被替代的领域。优秀的作品应是二者的结合：AI提供高效的数据基础设施，而设计师则赋予路线以体验性的意义结构。

从更广阔的历史视角来看，我们应认识到，一种工作能否被视为具有智识创造意义的“研究”，其标准始终在变化[10]。例如，描述性统计

报告在19世纪曾被视为严肃的学术贡献，而今天已很少有人如此认为。在景观设计领域，效果图制作也曾是一项重要而有价值的设计工种，如今却正在迅速被自动化工具所替代。表面上看，标准的提升似乎只是“水涨船高”的结果，而其内在动机则是随着知识前沿的推进和工具能力的增强，那些曾经位于前沿的工作，逐渐被纳入可内插的范畴——起初由人工完成，随后由机器。AI并未改变这一历史进程的方向，只是极大地加速了它。

由此，我们或许可以提出一个更一般性的准则：如果一项工作能够被定义为既有方法框架在新数据或新场景中的应用，那么无论它在当下看起来多么技术密集，其本质仍是内插性的。这并不是贬低。如前所述，内插性工作具有重要的实践价值，但其学术生命终究是有限的。相反，在既有知识尚无法自动生成答案之处，承担方向设定、价值判断与经验赋义的责任，才是外推性工作的应有之义。

那么，在景观设计的实践与研究语境中，什么才是真正的外推性工作？是发明全新的空间组织原则，是发现人与环境交互的未知机制，还是重新定义“何为好的景观”的价值命题？这些问题并不存在现成答案。但这恰是关键所在——提出一个真正好的问题，本身就是外推的纯粹形态。

对现有从业者和学生而言，如何转向外推性工作，自然也不会是轻而易举的。除了更高的智识要求之外，新的劳动分工对行业既有结构的冲击，或许构成了更现实的张力。仍以效果图制作为例，我们必须意识到，它不仅是一类设计工种，也是培养年轻设计师空间感知能力的重要训练路径。当这项工作被AI“解决”后，消失的不仅是一个岗位，更可能是一条通往设计直觉的学徒路径。这一挑战真实而深刻。或许，这本身也是时代留给我们的另一个外推性问题。

从长远来看，AI的能力无疑还将持续增长，也许会远远超出今天的想象。然而，即便如此——或者恰恰因此——在已知的前沿不断探索外推性的研究和设计，也许正构成了人类最独特、也最值得的智识承诺。正如科学共同体在80年前所写下的那样：前沿是无尽的[11]。在AI时代，我们不应仅将这句话视为一种修辞；它比以往任何时候都更成其为一个邀请。