

SpatialGuard: Harness-Guided Verifiable Spatial Reasoning for Text-to-Image Generation

Ziyun Qian^{1,2}, Zizhi Chen^{1,2}, Yizhou Liu^{1,2}, Mingyang Sun^{1,2},
Dingkang Yang^{1,2,*†}, Lihua Zhang^{1,2,*}

¹College of Intelligent Robotics and Advanced Manufacturing, Fudan University

²Fysics Intelligence Technologies Co., Ltd. (Fysics AI)

Abstract

Complex 3D spatial text to image generation requires models to convert natural language into stable visual geometry, not merely semantic appearance. Existing prompt-driven or layout-conditioned methods improve controllability, but often lack an optimizable and verifiable spatial intermediary before visual sampling. As a result, object relations, occlusion, visibility, and camera constraints can decay during multi-round generation. This paper presents SpatialGuard, a structured layout-guided framework for complex 3D spatial text-to-image generation. SpatialGuard parses prompts into image synthesis-oriented 3D layouts through a Spatial Layout Architect, realizes them as visual conditions and candidate images through a Visual Realizer, and uses a Visual Alignment Critic to validate consistency among prompt, layout, and image. To keep constraints stable across iterations, SpatialGuard introduces a Layout Harness that organizes rule constraints, tool invocation, shared knowledge, and feedback loops around the editable layout state. This design turns complex spatial generation from implicit prompt following into a verifiable process of planning, realization, validation, and repair. Comprehensive experiments show that SpatialGuard achieves state-of-the-art performance in complex 3D spatial layout generation and improves spatial faithfulness over existing text-to-image and layout control baselines.

1 Introduction

Complex 3D spatial text-to-image generation is a key capability for bridging natural-language interaction and controllable visual content production, with important applications in game-asset previewing (Husen et al., 2025; Li et al., 2025b), film storyboard design (Zhang et al., 2025a; Wang et al., 2026a), and embodied-intelligence simula-

tion (Feng et al., 2026). These applications require models to translate spatial intent in natural language into stable visual geometry, rather than merely generating semantically related appearance content. Traditional prompt-driven methods (Pan et al., 2026a; Qiu et al., 2025; Han et al., 2026a; Jiang et al., 2026a) mainly rely on text-to-image models to infer spatial structure from text implicitly. Thus, in relation-dense, view-sensitive, and layout-constrained scenes, they often satisfy only local semantics and struggle to maintain global spatial consistency. Recent methods (Venkatesh et al., 2026; Qin et al., 2026; Agrawal et al., 2026; Zhang et al., 2025b) introduce layout conditions, hierarchical generation, 3D control signals, and multi-step planning to make spatial structure more explicit during generation, significantly improving the controllability of complex scenes. However, existing methods (Parihar et al., 2025; Li et al., 2025a) still face two key limitations. First, the generation process usually lacks an image-synthesis-oriented spatial layout intermediary that is optimizable and verifiable. As a result, complex spatial intent lacks a stable carrier before visual sampling, and text alone fails to provide sufficiently explicit and executable spatial constraints for subsequent synthesis. Second, when complex spatial intent requires multiple rounds of planning, generation, evaluation, and correction, large models are prone to context forgetting and constraint decay. Spatial requirements established in early stages are difficult to preserve in later interactions, weakening fidelity to the original prompt and degrading generation quality.

These limitations highlight two key requirements for complex spatial text-to-image generation. First, the model needs to form a verifiable layout representation before image sampling, allowing spatial constraints in text to be explicitly encoded, continuously optimized, and checked before generation. Second, the system needs a cross-round

*Corresponding authors. †Project lead.




Figure 1: Example results of SpatialGuard on complex spatial text-to-image generation. Each case presents the input prompt, the generated 3D cube layout, and the corresponding synthesized image across urban street, farmyard, study, and bedroom scenes, demonstrating the model’s ability to preserve multi-object spatial relations, directional constraints, and scene-style descriptions.

structured constraint enforcement mechanism, so that spatial requirements do not remain merely in prompts or scoring signals but are stably preserved, actively retrieved, and updated according to feedback across planning, execution, verification, and revision. The former provides explicit spatial targets for image synthesis, while the latter ensures that these targets remain effective in long-chain interactions despite memory decay. Together, they indicate a stronger generation paradigm: complex 3D layouts should not only be prompted by language, but also explicitly designed, executed, and maintained.

To address these issues, we propose SpatialGuard, a structured-layout-guided generation framework for complex 3D spatial text-to-image generation. It unifies verifiable layout representation and persistent constraint enforcement within a single generation process. SpatialGuard consists of three collaborative modules. The Spatial Layout Architect parses input text into an image-synthesis-oriented 3D structured layout, providing a checkable and editable representation of complex spatial intent before sampling. The Visual Realizer converts this layout into visual conditions and synthesizes images, bringing spatial targets into pixel space. The Visual Alignment Critic evaluates con-

sistency among text, layout, and image, and feeds deviations back to subsequent planning. These modules operate around a shared layout state, preventing text understanding, image realization, and result evaluation from becoming decoupled across multi-round interaction.

Moreover, to prevent memory decay of spatial constraints during multi-round agent planning and interaction, we first introduce the idea of an agent harness (Lin et al., 2026; Li et al.; Pan et al., 2026b) into text-to-image generation and distill it into four core mechanisms: rule constraints, tool invocation, shared knowledge, and feedback loops. Rule constraints convert spatial requirements in language into checkable layout boundaries. Tool invocation turns layout decisions into externally executable operations. Shared knowledge preserves layout versions, evaluation judgments, and revision trajectories across stages. Feedback loops reinject detected spatial deviations into layout planning and image generation. Based on the automated project workflow, SpatialGuard generates an initial 3D layout from text, performs deterministic verification according to the object list and spatial relations, iteratively invokes 3D layout rendering and image synthesis, and produces executable repair plans when spatial violations are detected. Thus,

it shifts complex spatial text-to-image generation from prompt-driven implicit generation to verifiable generation guided by structured layouts and maintained by constraint mechanisms.

The main contributions of this paper are as follows:

- We propose SpatialGuard, an agentic layout generation framework for complex 3D spatial text-to-image generation, enabling closed-loop modeling from textual spatial intent to verifiable visual generation.
- We introduce the concept of agent harness into the text-to-image pipeline for the first time and build a cross-round spatial constraint enforcement mechanism to mitigate memory decay in complex spatial generation by large models.
- Comprehensive quantitative and qualitative experiments show that SpatialGuard achieves new state-of-the-art performance in complex 3D spatial text-to-image generation and significantly outperforms existing text-to-image and layout-control baselines.

2 Related Work

2.1 Text to Image Generation

Text-to-image generation has progressed from high-fidelity synthesis to more deliberate modeling of semantic and spatial intent. General models improve prompt adherence through reasoning, reward optimization, autoregressive modeling, and spatial evaluation signals (Jiang et al., 2025; Han et al., 2026b; Jiang et al., 2026b,c,d). Janus-Pro-R1 (Pan et al., 2026a) and T2I-R1 (Jiang et al., 2026a) connect visual comprehension with generation through reinforcement learning and chain of thought style planning, while NextStep-1 (Han et al., 2026a) explores scalable autoregressive generation with continuous visual tokens (Han et al., 2026a). Self-Cross (Qiu et al., 2025) reduces subject mixing among similar objects, and SpatialScore provides a reward model for spatial relation evaluation and reinforcement learning (Tang et al., 2026). Recent benchmarks show that compositional and spatial alignment have become central axes for evaluating text-to-image models (Wang et al., 2026b,c; Huang et al., 2025). These works strengthen semantic alignment, but spatial composition is still largely learned or rewarded through

implicit model behavior. A complementary line of work makes spatial control more explicit through planning, layout, pose, and 3D conditions. LayoutGPT (Feng et al., 2023) uses language models for compositional visual planning and translates text requests into explicit scene programs that specify object categories and bounding boxes before rendering. MCCD (Li et al., 2025a) coordinates prompt level decomposition, image generation, and selection, whereas CREA (Venkatesh et al., 2026) assigns distinct creative and critical roles to refine visual concepts. LayerCraft (Zhang et al., 2025b) decomposes generation into layered reasoning and object integration, deriving a dependency aware scene graph that schedules background construction before foreground integration. Compass Control (Parihar et al., 2025) binds orientation to individual objects, SceneDesigner (Qin et al., 2026) encodes location, scale, and orientation in camera view pose maps, and SeeThrough3D (Agrawal et al., 2026) renders translucent cuboids that reveal occlusion and camera view during synthesis. Together, these interfaces make selected spatial variables visible to the generator, although their intermediate representations remain primarily fixed conditions rather than states designed for iterative verification and repair. These studies show that layouts, poses, layers, and planning improve controllability. However, existing methods still lack an optimizable and verifiable spatial layout intermediary, leaving complex spatial intent without a stable carrier before visual sampling and weakening executable spatial constraints for synthesis.

2.2 Harness for Agentic Execution

Agent harness research studies the execution layer around a foundation model, covering tools, state, orchestration, validation, observability, and recovery. Recent surveys organize harness design around execution environment, tool interface, context management, lifecycle control, verification, and governance (Li et al.; Meng et al., 2026). Such decompositions separate execution concerns and attribute failures to a specified layer, rather than treating an agent run as one opaque model call. Agentic Harness Engineering evolves harness components through structured observability over files, trajectories, and decisions, exposes reusable components as files, distills layered execution traces into evidence, and tests every revision against a declared prediction (Lin et al., 2026). Natural Lan-




Figure 2: Overview of SpatialGuard. Given an input prompt, the Spatial Layout Architect parses objects, spatial relations, and scene appearance into a verifiable 3D layout under the Layout Harness. The Visual Realizer converts the layout into masks, token-aligned conditions, and a candidate image. The Visual Alignment Critic performs structured validation over object placement, relations, visibility, and camera constraints. Failed constraints are converted into executable repair actions and fed back to update the layout, while rule constraints, tool invocation, shared knowledge, and feedback loops preserve spatial intent across iterations until the final image satisfies the prompt.

guage Agent Harnesses separate readable harness policy from runtime mechanisms; their documents specify run policy while the shared runtime executes calls, state updates, and validation gates (Pan et al., 2026b). Meta Harness treats harness optimization as an outer loop over code, traces, and scores, allowing a proposer to inspect prior code, trajectories, and scores through a persistent filesystem (Lee et al., 2026). Beyond coding agents, harness-style designs have begun to appear in visual and embodied tasks. VASA maintains a persistent working mask for open ad hoc segmentation, turning visual progress into an inspectable state that persists across interactions (Wang and Yu, 2026). SceneWeaver uses a reflective tool-based loop to synthesize 3D scenes under semantic and physical feedback, selecting modular scene synthesis tools and revising scenes after each evaluation (Yang et al., 2026). These systems indicate that explicit state, callable tools, validation gates, and shared memory can help complex layout text-to-image generation keep spatial intent readable,

enforceable, and recoverable.

3 Methodology

3.1 Overview

Complex 3D spatial text-to-image generation requires preserving object identity, relative position, visibility, camera configuration, and realism from language understanding to image synthesis. Existing controllable methods use layouts, 3D controls, layered generation, or agentic planning to improve spatial grounding (Agrawal et al., 2026; Qin et al., 2026; Parihar et al., 2025; Zhang et al., 2025b; Li et al., 2025a), but spatial requirements often remain scattered across prompts, conditions, and evaluation signals. Relations understood early may weaken after synthesis or disappear during revision. SpatialGuard addresses this by combining agent-based spatial reasoning with a Layout Harness, as shown in Fig. 2.

Given an input prompt, SpatialGuard separates object mentions, spatial intentions, and appearance

descriptions, then organizes them into a layout-centered generation loop. The system plans a structured scene, realizes it as visual evidence, checks whether the image follows the spatial intent, and repairs the layout when violations appear. In Fig. 2, the input enters the Spatial Layout Architect, the Visual Realizer produces a candidate image, the Visual Alignment Critic performs structured validation, and failed items return through an executable repair plan.

The Layout Harness keeps this loop stable. Rule constraints make object, relation, scale, support, occlusion, visibility, and camera requirements checkable. Tool invocation maps revisions to operations such as moving objects, rescaling objects, and adjusting the camera. Shared knowledge stores layout decisions, validation records, and repair traces across rounds. Feedback loops send detected spatial errors back to planning rather than leaving them as passive scores. SpatialGuard therefore turns prompt-only sampling into a verifiable process that repeatedly plans, realizes, validates, and repairs the scene.

3.2 SpatialGuard Framework

SpatialGuard instantiates the common interface in Sec. 3.1 as an agentic layout generation framework. Prior layout-driven and 3D-controlled text-to-image methods make spatial control more explicit (Feng et al., 2023; Agrawal et al., 2026; Qin et al., 2026; Parihar et al., 2025), but often pass the layout as a static generator condition. SpatialGuard instead treats the layout as an editable state that links language understanding, visual synthesis, and result validation, moving complex 3D spatial intent from implicit prompt attention into a state that can be planned, grounded, and inspected.

Following Sec. 3.1, p denotes the input prompt, $\mathcal{O} = \{\mathbf{o}_i\}_{i=1}^N$ denotes the set of N mentioned objects, $\mathbf{o}_i$ denotes the i th object. $\mathcal{C}$ denotes the spatial constraint set extracted from p , and $\mathbf{a}$ denotes the appearance and environment description. At round t , the layout state is written as

$$\mathbf{L}^{(t)} = \left(\left\{ \left(\mathbf{c}_i, \mathbf{x}_i^{(t)}, \mathbf{s}_i^{(t)}, \boldsymbol{\alpha}_i^{(t)} \right) \right\}_{i=1}^N, \boldsymbol{\kappa}^{(t)} \right), \quad (1)$$

where $\mathbf{c}_i$ is the category of object $\mathbf{o}_i$, $\mathbf{x}_i^{(t)}$ is the 3D position of object $\mathbf{o}_i$ at round t , $\mathbf{s}_i^{(t)}$ is its layout-readable size at round t , $\boldsymbol{\alpha}_i^{(t)}$ is its azimuth at round t , and $\boldsymbol{\kappa}^{(t)}$ is the camera parameter vector at round

t . SpatialGuard operates on this state through three cooperative modules.

The **Spatial Layout Architect** converts the parsed spatial intent into an executable geometry carrier. It initializes the first layout and later revises the layout according to structured feedback:

$$\begin{aligned} \mathbf{L}^{(0)} &= \mathcal{A}_s(p, \mathcal{O}, \mathcal{C}, \mathbf{a}), \\ \mathbf{L}^{(t+1)} &= \mathcal{A}_s(p, \mathcal{O}, \mathcal{C}, \mathbf{a}, \mathbf{L}^{(t)}, \boldsymbol{\Delta}^{(t)}), \end{aligned} \quad (2)$$

where $\mathcal{A}_s$ denotes the Spatial Layout Architect, $\mathbf{L}^{(0)}$ denotes the initial layout state, $\mathbf{L}^{(t+1)}$ denotes the next layout state after revision, and $\boldsymbol{\Delta}^{(t)}$ denotes the repair feedback produced by the Visual Alignment Critic at round t . The architect resolves object mentions, normalizes categories and counts, decomposes spatial phrases into 3D and screen space requirements, and assigns positions, sizes, orientations, and camera parameters. For example, a front right relation is represented through depth ordering and lateral placement, keeping it checkable after projection.

The **Visual Realizer** maps the current layout state into visual conditions and synthesizes a candidate image. Given $\mathbf{L}^{(t)}$, it renders the layout, derives instance-level masks, aligns object regions with the corresponding prompt tokens, and invokes the generator:

$$\mathbf{I}^{(t)} = \mathcal{G}\left(p, \mathbf{a}, \rho\left(\mathbf{L}^{(t)}\right), \psi\left(p, \mathbf{L}^{(t)}\right)\right), \quad (3)$$

where ρ denotes the layout renderer, ψ denotes the operator that produces instance masks and token-aligned object regions, $\mathcal{G}$ denotes the layout-conditioned image generator, and $\mathbf{I}^{(t)}$ denotes the candidate image generated in round t . The rendered layout provides global geometry, masks preserve object boundaries, and token alignment keeps appearance phrases attached to targets, transferring the planned 3D arrangement into pixel space.

The **Visual Alignment Critic** closes the loop by judging whether the candidate image respects the prompt and the current layout:

$$\left(\mathbf{v}^{(t)}, \boldsymbol{\Delta}^{(t)}\right) = \mathcal{A}_c\left(p, \mathcal{C}, \mathbf{L}^{(t)}, \mathbf{I}^{(t)}\right), \quad (4)$$

where $\mathcal{A}_c$ denotes the Visual Alignment Critic, $\mathbf{v}^{(t)}$ is a structured validation record, and $\boldsymbol{\Delta}^{(t)}$ is the repair feedback for the next layout revision. The critic reads $\mathcal{C}$, which contains the required object relations, support conditions, visibility requirements, and camera-related constraints extracted

from $\mathbf{p}$. It localizes failures to concrete fields of $\mathbf{L}^{(t)}$, such as object position, size, azimuth, occlusion status, or camera configuration, instead of collapsing the judgment into a single scalar score. If $\mathbf{v}^{(t)}$ contains no failed item, SpatialGuard returns $\mathbf{I}^{(t)}$ as the final output. Otherwise, $\Delta^{(t)}$ is passed back to the Spatial Layout Architect to produce $\mathbf{L}^{(t+1)}$.

This separation gives each module a clear responsibility: the Spatial Layout Architect makes spatial intent explicit, the Visual Realizer turns it into visual evidence, and the Visual Alignment Critic converts deviations into actionable layout feedback. SpatialGuard, therefore, supports closed-loop modeling from textual spatial intent to verifiable visual generation without requiring one-pass implicit spatial inference by the image generator.

3.3 Layout Harness for Persistent Spatial Execution

Existing text-to-image systems make spatial control more explicit through layouts, 3D conditions, layered composition, and agent collaboration (Feng et al., 2023; Agrawal et al., 2026; Qin et al., 2026; Parihar et al., 2025; Zhang et al., 2025b; Li et al., 2025a). However, these signals are often used as local prompts, static conditions, or post hoc scores, which are fragile for relation-dense scenes. A relation parsed early may vanish during object repair, visibility adjustment, or camera change. The core challenge is therefore to preserve, retrieve, and enforce relations across planning, realization, validation, and revision.

Agent harness research provides an execution view for this challenge (Lin et al., 2026; Li et al.; Pan et al., 2026b). Although a harness has no single mathematical definition, prior work commonly treats it as an external layer that organizes model calls, tools, state, validation, recovery, and stopping conditions. From this view, we distill four properties for complex spatial generation: rule constraints, tool invocation, shared knowledge, and feedback loops. SpatialGuard instantiates them as a Layout Harness around the editable layout state introduced in Sec. 3.2.

At refinement round t , where t indexes the current iteration, the Layout Harness maintains

$$\mathbf{H}^{(t)} = (\mathcal{C}, \mathcal{T}, \mathbf{M}^{(t)}). \quad (5)$$

Here $\mathbf{H}^{(t)}$ is the harness state at round t , $\mathcal{C}$ is

the spatial constraint set extracted from the input prompt $\mathbf{p}$, $\mathcal{T}$ is the executable layout tool library, and $\mathbf{M}^{(t)}$ is the shared knowledge state. $\mathbf{M}^{(t)}$ records the object manifest, parsed spatial predicates, previous layout states, validation records, and repair actions, giving all modules a durable reference beyond a single model call.

Rule constraints make linguistic requirements inspectable. SpatialGuard stores relations such as front, left, support, or full visibility in $\mathcal{C}$ and checks them against $\mathbf{L}^{(t)}$ and $\mathbf{I}^{(t)}$. As defined in Sec. 3.2, $\mathbf{L}^{(t)}$ is the 3D layout state at round t , and $\mathbf{I}^{(t)}$ is the candidate image generated at round t . The Visual Alignment Critic produces $\mathbf{v}^{(t)}$ and $\Delta^{(t)}$, where $\mathbf{v}^{(t)}$ lists passed and failed constraints and $\Delta^{(t)}$ specifies the next layout revision.

Tool invocation turns a revision instruction into an executable operation. Given the validation result, the harness selects and applies a sequence of tool actions:

$$\begin{aligned} \mathbf{u}^{(t)} &= \Gamma(\mathbf{v}^{(t)}, \Delta^{(t)}, \mathbf{H}^{(t)}), \\ \tilde{\mathbf{L}}^{(t+1)} &= \Omega(\mathbf{L}^{(t)}, \mathbf{u}^{(t)}), \\ \mathbf{M}^{(t+1)} &= \Xi(\mathbf{M}^{(t)}, \tilde{\mathbf{L}}^{(t+1)}, \mathbf{v}^{(t)}, \mathbf{u}^{(t)}). \end{aligned} \quad (6)$$

In this equation, $\mathbf{u}^{(t)}$ is the selected sequence of tool actions at round t , Γ is the tool selection function, Ω is the execution function that applies $\mathbf{u}^{(t)}$ to $\mathbf{L}^{(t)}$, $\tilde{\mathbf{L}}^{(t+1)}$ is the tool-executed intermediate layout state, Ξ is the memory update function, and $\mathbf{M}^{(t+1)}$ is the updated shared knowledge state. The tools include object translation, object scaling, support adjustment, occlusion control, visibility restoration, focal length adjustment, and camera elevation adjustment. Because they operate on explicit layout fields, a failed relation can be repaired while preserving constraints that already pass validation.

Shared knowledge keeps spatial decisions recoverable across rounds. The state $\mathbf{M}^{(t)}$ tells later modules which objects exist, which relations come from $\mathbf{p}$, which fields of $\mathbf{L}^{(t)}$ changed, and which failures remain. This avoids reliance on incomplete conversation history and makes each image traceable to a layout version, validation record, and repair action sequence.

Feedback loops connect validation to the next execution step. When $\mathbf{v}^{(t)}$ contains failed constraints, $\Delta^{(t)}$ is routed through Γ to obtain $\mathbf{u}^{(t)}$, and Ω produces the tool-executed intermediate layout state

Table 1: Quantitative comparison of spatial layout faithfulness judged by three vision language models (Google DeepMind, 2025; Qwen Team, 2025; xAI, 2025). Scores range from 1 to 10 and may be fractional, with all values reported to two decimal places. Higher scores are better. The best results are highlighted in bold and the second best results are underlined.

Method	Presence $\uparrow$	Position $\uparrow$	Relation $\uparrow$	Depth $\uparrow$	Scale $\uparrow$	Support $\uparrow$	Framing $\uparrow$	Overall $\uparrow$
HunyuanImage-2.1 (Tencent, 2025)	<u>8.59</u>	7.98	7.84	7.25	7.79	<u>7.64</u>	8.18	<u>7.90</u>
SD-3.5-L (Esser et al., 2024)	8.31	<u>8.08</u>	7.70	7.34	<u>7.91</u>	7.37	<u>8.46</u>	7.88
FLUX.1-dev (Black Forest Labs, 2024)	8.13	7.62	<u>7.96</u>	<u>7.49</u>	7.57	7.18	7.92	7.70
Self-Cross (Qiu et al., 2025)	6.06	5.20	4.96	4.74	4.87	4.82	6.80	5.35
T2I-R1 (Jiang et al., 2026a)	6.49	5.91	6.62	6.25	5.58	5.39	6.68	6.13
Janus-Pro-R1 (Pan et al., 2026a)	6.83	6.25	6.55	6.18	6.04	5.66	6.55	6.29
NextStep-1 (Han et al., 2026a)	7.22	6.72	6.35	6.05	6.36	6.08	6.95	6.53
SpatialGuard (Ours)	9.26	8.83	9.38	9.67	9.39	9.45	9.58	9.37

$\tilde{\mathbf{L}}^{(t+1)}$, which is then used by the Spatial Layout Architect to form the next layout state $\mathbf{L}^{(t+1)}$. When no failed constraint remains, SpatialGuard returns $\mathbf{I}^{(t)}$ as the final image. Through this harness, text to image generation becomes a persistent spatial execution process in which language constraints are externalized, checked, repaired, and remembered until the layout and image satisfy the requested spatial intent.

4 Experiments

4.1 Implementation Details

SpatialGuard is a training-free framework, and all model weights remain frozen during evaluation. We conduct all experiments on a single NVIDIA H200 GPU with 143 GB of memory. Spatial Layout Architect and Visual Alignment Critic are implemented with GPT5 (OpenAI, 2025), where the former produces executable 3D layouts and the latter returns structured validation records and repair instructions. Visual Realizer is built on FLUX.1 dev (Black Forest Labs, 2024). We render layouts at 1024 resolution and use 512 resolution conditions for image synthesis. The default inference setting uses 25 denoising steps, a guidance scale of 3.5, and up to 4 layout verification and repair rounds.

4.2 Quantitative Evaluation

We evaluate spatial layout faithfulness with Gemini 2.5 Pro (Google DeepMind, 2025), Qwen3-VL (Qwen Team, 2025), and Grok 4 (xAI, 2025) using a fixed prompt set shared by all methods. The prompts cover object presence, directional relations, depth ordering, support, relative scale, and camera framing. We compare SpatialGuard with representative text to image methods (Tencent, 2025; Esser et al., 2024; Qiu et al., 2025).

For each prompt, every method produces three final images with its recommended settings. Evaluators receive only the image and its prompt, without method identities or intermediate artifacts such as layout renderings, masks, validation records, or repair traces. This blind protocol uses identical prompts and instructions for all systems. No method is tuned on the evaluation prompts, and every comparison uses the same number of images. All three models score each image once, and the reported score for each dimension averages their judgments over all prompts and generated images.

Each image is evaluated on seven dimensions. Presence measures whether all requested objects and counts are visible and separable, with penalties for omissions, count errors, or severe fusion. Position measures how closely object locations and requested anchors match the prompt, while allowing partial credit for approximate placement. Relation evaluates atomic and compound spatial relations, with partial credit when only some components are satisfied. Depth covers front and behind order, occlusion, and residual visibility. Scale measures relative object sizes, distances between objects, and camera distance. Support evaluates contact with the specified surface and physical plausibility. Framing measures how well viewpoint, cropping, object visibility, and composition expose the requested layout. Overall is the arithmetic mean of these seven dimension scores. Appearance quality, style, and photorealism are ignored unless they prevent a reliable spatial judgment.

Evaluators may assign decimal scores on a 1 to 10 scale, and all reported values are rounded to two decimal places. Scores from 9 to 10 indicate clear and nearly complete satisfaction; 7 to 8 indicate a minor deviation; 5 to 6 indicate partial satisfaction or substantial ambiguity; 3 to 4 indicate major violations while relevant objects remain recog-

Input Prompt: A sofa anchors the image. A cat lounges clearly front-left of the sofa, while a dog rests clearly front-right of the sofa. A lamp glows to the left of the sofa, and a bookshelf on the right of the sofa. The scene is a quiet and comfortable pet resting corner, soft warm indoor lighting, lazy and relaxed home atmosphere, naturally arranged objects, rich details, and a warm realistic composition.



Input Prompt: A bus centers the image. A man riding a bicycle is front-left of the bus. A sedan is located to the right of the bus. A bird is flying back-right of the bus. The scene is a quiet city curbside near a bus stop, with realistic daylight, street-side details, and a clean documentary travel style.



Figure 3: Qualitative comparison on complex spatial prompts. SpatialGuard preserves object completeness and relative layout more reliably than baselines, which often miss objects, truncate visible entities, or violate directional relations.

nizable; and 1 to 2 indicate missing, contradictory, severely occluded, or unjudgeable evidence. Compound relations receive partial credit for visibly satisfied components. If a required object is absent, the associated relation receives a score from 1 to 2. Evaluators use only visible evidence and assign the lower score when evidence is ambiguous.

Table 1 shows that SpatialGuard obtains the highest Overall score of 9.37 and ranks first on all seven dimensions. HunyuanImage-2.1 (Tencent, 2025) is the strongest baseline at 7.90, narrowly ahead of SD-3.5-L (Esser et al., 2024) at 7.88, leaving a 1.47 point gap to SpatialGuard. The largest gains over the best baseline for each dimension appear on Depth, Support, Scale, and Relation, at 2.18, 1.81, 1.48, and 1.42 points. Presence improves by 0.67, indicating that the main benefit lies in spatial structure rather than simple object inclusion. The Spatial Layout Architect converts language into explicit positions, depth order, orientations, and camera parameters, while the Visual Realizer transfers this plan into pixel space. The Visual Alignment Critic identifies visible violations, and the Layout Harness preserves satisfied constraints during targeted repair. Their coordination reduces relation reversals, inconsistent occlusion, and physically implausible layouts, which accounts for the consistent advantage across the spatial metrics.

The smaller gain on Presence clarifies where the method helps most. Strong generators usually recover common objects, but they remain less reliable when several constraints must hold simultaneously. SpatialGuard instead preserves those constraints across planning, synthesis, inspection, and correction, yielding balanced performance.

4.3 Qualitative Evaluation

Figure 3 compares SpatialGuard with recent text-to-image baselines on complex prompts containing object completeness, directional relations, and scene-level constraints. Baselines often produce visually plausible images, but their spatial structure is unstable. In the sofa scene, the prompt requires the cat and dog to appear in front of the sofa, while Self cross fails to preserve this layout. FLUX.1 dev (Black Forest Labs, 2024) generates only a partial bookshelf, and T2I-R1 (Jiang et al., 2026a) omits the lamp and bookshelf in the first example. In the bus scene, Janus-pro-r1 (Pan et al., 2026a) misses the bird, breaking the requested back right relation. SpatialGuard avoids these failures by maintaining an explicit layout state and checking object, relation, and visibility constraints before accepting the output. This leads to images that are not only realistic but also faithful to the full spatial intent.

Across both examples, the key distinction is not

Table 2: Ablation study using the same three vision language model evaluators (Google DeepMind, 2025; Qwen Team, 2025; xAI, 2025) and metrics reported in Table 1. Scores range from 1 to 10, and higher scores are better. The best results are highlighted in bold and the second best results are underlined.

Method	Presence $\uparrow$	Position $\uparrow$	Relation $\uparrow$	Depth $\uparrow$	Scale $\uparrow$	Support $\uparrow$	Framing $\uparrow$	Overall $\uparrow$
Ours w/o Spatial Layout Architect	<u>9.01</u>	8.10	8.62	8.70	8.52	8.40	9.00	8.62
Ours w/o Visual Realizer	8.74	<u>8.55</u>	8.86	9.03	8.76	<u>8.95</u>	8.84	8.82
Ours w/o Visual Alignment Critic	8.99	8.38	<u>9.08</u>	<u>9.32</u>	9.02	8.84	<u>9.31</u>	8.99
Ours w/o harness	8.92	8.47	9.00	9.24	<u>9.18</u>	9.03	9.22	<u>9.01</u>
Ours(full)	9.26	8.83	9.38	9.67	9.39	9.45	9.58	9.37

isolated visual quality but whether every requested constraint remains usable within the same composition. A plausible image can still fail the scene specification when an object is absent, a directional relation is reversed, or view dependent placement obscures an entity needed by another relation. SpatialGuard makes these deviations explicit at the layout level, where object identity, relative position, and visibility are checked before an image is accepted. This prevents a correction to one element from silently disrupting another and explains why the method preserves complete, readable, and spatially coherent scenes after rendering. The missing bookshelf and lamp, the lost bird, and the failed foreground relation show distinct ways prompt information can disappear during synthesis. The layout state carries these requirements forward as inspectable commitments, and validation identifies each item requiring specific repair.

4.4 Ablation Studies

Table 2 evaluates the four components with the same vision language model protocol (Google DeepMind, 2025; Qwen Team, 2025; xAI, 2025) as Table 1. The complete system reaches an Overall score of 9.37. The second best score changes by criterion: the Architect ablation leads Presence at 9.01, the Realizer ablation leads Position and Support at 8.55 and 8.95, the Critic ablation leads Relation, Depth, and Framing at 9.08, 9.32, and 9.31, and the Harness ablation leads Overall at 9.01. This distribution follows the roles of the components. Without the Architect, Presence remains high, but Position and Relation fall to 8.10 and 8.62, showing that precise geometry needs an explicit layout plan. The Realizer transfers this plan into image evidence, while the Critic detects visible violations for repair. The Harness preserves decisions across rounds, affecting Scale, Support, and Overall. The gap from the full configuration shows that SpatialGuard benefits from their interaction rather than one isolated module.

These patterns clarify the results. Removing the Architect preserves basic object presence but produces the largest loss in position and relation, indicating that the layout state supplies one common reference for subsequent stages. The Realizer grounds planned geometry in the image, while the Critic turns visible discrepancies into repair signals. The Harness ablation shows that persistent constraints support SpatialGuard when planning, realization, and revision must remain coordinated across rounds.

5 Conclusion

This paper introduces SpatialGuard, a structured layout-guided framework for complex 3D spatial text-to-image generation. SpatialGuard externalizes spatial intent into an editable 3D layout state and organizes generation as planning, realization, validation, and repair, rather than leaving spatial relations to implicit prompt interpretation. Its Layout Harness stabilizes this process with rule constraints, tool invocation, shared knowledge, and feedback loops, keeping object relations, visibility, and camera constraints recoverable across iterations. By treating the layout as a shared state, the framework keeps language, visual synthesis, and evaluation anchored to shared spatial decisions. Corrective actions remain interpretable because a detected violation maps to an object, relation, or camera choice instead of an opaque sampling trajectory. Experiments show clear gains in spatial pose, relation grounding, and measurement accuracy over strong text-to-image baselines. Qualitative and ablation results further show these gains come from verifiable layout execution.

6 Acknowledgements

This work is supported by the 2026 AI Breakthrough Initiative Research Project 2026JLGJ0001GX.

Limitations

SpatialGuard focuses on spatially grounded scenes that can be expressed through object lists, relations, visibility constraints, and camera parameters. Prompts with highly abstract artistic intent or intentionally ambiguous spatial descriptions may require additional interpretation rules. Since the framework performs planning, validation, and repair before producing the final image, its inference cost is higher than a single direct call to a text-to-image model. In addition, the final visual quality still depends on the capability of the underlying image generator. Future work can extend the tool library to finer physical interactions and improve efficiency through lighter validation strategies.

References

- Vaibhav Agrawal, Rishubh Parihar, Pradhaan S Bhat, Ravi Kiran Sarvadevabhatla, and Venkatesh Babu Radhakrishnan. 2026. Seethrough3d: Occlusion aware 3d control in text-to-image generation.
- Black Forest Labs. 2024. FLUX. <https://github.com/black-forest-labs/flux>. Accessed: 2026-05-26.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, and 1 others. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Weixi Feng, Wanrong Zhu, Tsu-jui Fu, Varun Jampani, Arjun Akula, Xuehai He, Sugato Basu, Xin Eric Wang, and William Yang Wang. 2023. Layoutgpt: Compositional visual planning and generation with large language models. *Advances in Neural Information Processing Systems*, 36:18225–18250.
- Zhaohan Feng, Ruiqi Xue, Lei Yuan, Yang Yu, Ning Ding, Meiqin Liu, Bingzhao Gao, Jian Sun, Xinhui Zheng, and Gang Wang. 2026. Multi-agent embodied ai: Advances and future directions. *Science China Information Sciences*, 69(5):151202.
- Google DeepMind. 2025. Gemini 2.5: Our most intelligent AI model. <https://blog.google/innovation-and-ai/models-and-research/google-deepmind/gemini-model-thinking-updates-march-2025/>.
- Chunrui Han, Guopeng Li, Jingwei Wu, Quan Sun, Yan Cai, Yang Peng, Zheng Ge, Deyu Zhou, Haomiao Tang, Hongyu Zhou, and 1 others. 2026a. Nextstep-1: Toward autoregressive image generation with continuous tokens at scale. In *The Fourteenth International Conference on Learning Representations*.
- Minghao Han, Dingkan Yang, Yue Jiang, Yizhou Liu, and Lihua Zhang. 2026b. Omnifysics: Towards physical intelligence evolution via omni-modal signal processing and network optimization. *arXiv preprint arXiv:2602.07064*.
- Kaiyi Huang, Chengqi Duan, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. 2025. T2i-compbench++: An enhanced and comprehensive benchmark for compositional text-to-image generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(5):3563–3579.
- Dharma Hutama Husen, Wirawan Istiono, and 1 others. 2025. Procedural story generation for visual novels using large language models and text-to-image techniques. *Journal of Games, Game Art, and Gamification*, 10(3):106–114.
- Dongzhi Jiang, Ziyu Guo, Renrui Zhang, Zhuofan Zong, Hao Li, Le Zhuo, Shilin Yan, Pheng-Ann Heng, and Hongsheng Li. 2026a. T2i-r1: Reinforcing image generation with collaborative semantic-level and token-level cot. *Advances in Neural Information Processing Systems*, 38:39856–39890.
- Yue Jiang, Xue Jiang, Lihua Zhang, Zhiqiang Wang, Yuhang Lu, Peng Wang, Bo Han, Feng Zheng, and Dingkan Yang. 2026b. Mm-snowball: Evaluating and mitigating hallucination snowballing in multimodal multi-turn dialogue. *arXiv preprint arXiv:2606.00622*.
- Yue Jiang, Jichu Li, Yang Liu, Dingkan Yang, Feng Zhou, and Quyu Kong. 2026c. Danmakutppbench: A multi-modal benchmark for temporal point process modeling and understanding. *Advances in Neural Information Processing Systems*, 38.
- Yue Jiang, Haiwei Xue, Minghao Han, Mingcheng Li, Xiaolu Hou, Dingkan Yang, Lihua Zhang, and Xu Zheng. 2026d. Satiredecoder: Visual cascaded decoupling for enhancing satirical image comprehension. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pages 5468–5476.
- Yue Jiang, Dingkan Yang, Minghao Han, Jinghang Han, Zizhi Chen, Yizhou Liu, Mingcheng Li, Peng Zhai, and Lihua Zhang. 2025. Fysicsworld: A unified full-modality benchmark for any-to-any understanding, generation, and reasoning. *arXiv preprint arXiv:2512.12756*.
- Yoonho Lee, Roshen Nair, Qizheng Zhang, Kangwook Lee, Omar Khattab, and Chelsea Finn. 2026. Meta-harness: End-to-end optimization of model harnesses. *arXiv preprint arXiv:2603.28052*.
- Junjie Li, Xi Xiao, Yunbei Zhang, Chen Liu, Lin Zhao, Xiaoying Liao, Yingrui Ji, Janet Wang, Jianyang Gu, Yingqiang Ge, and 1 others. Agent harness engineering: A survey.
- Mingcheng Li, Xiaolu Hou, Ziyang Liu, Dingkan Yang, Ziyun Qian, Jiawei Chen, Jinjie Wei, Yue Jiang, Qingyao Xu, and Lihua Zhang. 2025a. Mccd:

- Multi-agent collaboration-based compositional diffusion for complex text-to-image generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 13263–13272.
- Ruihuang Li, Caijin Zhou, Shoujian Zheng, Jianxiang Lu, Jiabin Huang, Comi Chen, Junshu Tang, Guangzheng Xu, Jiale Tao, Hongmei Wang, and 1 others. 2025b. Hunyuan-game: Industrial-grade intelligent game creation model. *arXiv preprint arXiv:2505.14135*.
- Jiahang Lin, Shichun Liu, Chengjun Pan, Lizhi Lin, Shihan Dou, Xuanjing Huang, Hang Yan, Zhenhua Han, and Tao Gui. 2026. Agentic harness engineering: Observability-driven automatic evolution of coding-agent harnesses. *arXiv preprint arXiv:2604.25850*.
- Qianyu Meng, Yanan Wang, Liyi Chen, Qimeng Wang, Chengqiang Lu, Wei Wu, Yan Gao, Yi Wu, and Yao Hu. 2026. Agent harness for large language model agents: A survey.
- OpenAI. 2025. Gpt-5 system card. <https://openai.com/index/gpt-5-system-card/>. Accessed: 2026-05-26.
- Kaihang Pan, Yang Wu, Wendong Bu, Shen Kai, Juncheng Li, Yingting Wang, Yunfei Li, Siliang Tang, Jun Xiao, Fei Wu, and 1 others. 2026a. Janus-pro-r1: Advancing collaborative visual comprehension and generation via reinforcement learning. *Advances in Neural Information Processing Systems*, 38:60013–60041.
- Linyue Pan, Lexiao Zou, Shuo Guo, Jingchen Ni, and Hai-Tao Zheng. 2026b. Natural-language agent harnesses. *arXiv preprint arXiv:2603.25723*.
- Rishubh Parihar, Vaibhav Agrawal, Sachidanand VS, and Venkatesh Babu Radhakrishnan. 2025. Compass control: Multi object orientation control for text-to-image generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 2791–2801.
- Zhenyuan Qin, Xincheng Shuai, and Henghui Ding. 2026. Scenedesigner: Controllable multi-object image generation with 9-dof pose manipulation. *Advances in Neural Information Processing Systems*, 38:133376–133400.
- Weimin Qiu, Jieke Wang, and Meng Tang. 2025. Self-cross diffusion guidance for text-to-image synthesis of similar subjects. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 23528–23538.
- Qwen Team. 2025. *Qwen3-VL technical report*. *arXiv preprint arXiv:2511.21631*.
- Zhenyu Tang, Chaoran Feng, Yufan Deng, Jie Wu, Xiaojie Li, Rui Wang, Yunpeng Chen, and Daquan Zhou. 2026. Enhancing spatial understanding in image generation via reward modeling. *arXiv preprint arXiv:2602.24233*.
- Tencent. 2025. HunyuanImage-2.1. <https://huggingface.co/tencent/HunyuanImage-2.1>. Accessed: 2026-05-26.
- Kavana Venkatesh, Connor Dunlop, and Pinar Yarnadag. 2026. Crea: A collaborative multi-agent framework for creative image editing and generation. *Advances in Neural Information Processing Systems*, 38:171332–171392.
- Xinran Wang, Songyu Xu, Shan Xiangxuan, Yuxuan Zhang, Muxi Diao, Xueyan Duan, Kongming Liang, Zhanyu Ma, and 1 others. 2026a. Cinetechbench: A benchmark for cinematographic technique understanding and generation. *Advances in Neural Information Processing Systems*, 38.
- Zehan Wang, Jiayang Xu, Ziang Zhang, Tianyu Pang, Chao Du, Hengshuang Zhao, and Zhou Zhao. 2026b. Genspace: Benchmarking spatially-aware image generation. *Advances in Neural Information Processing Systems*, 38.
- Zengbin Wang, Xuecai Hu, Yong Wang, Feng Xiong, Man Zhang, and Xiangxiang Chu. 2026c. Everything in its place: Benchmarking spatial intelligence of text-to-image models. *arXiv preprint arXiv:2601.20354*.
- Zilin Wang and Stella X Yu. 2026. Vision harnessing agent for open ad-hoc segmentation. *arXiv preprint arXiv:2605.19410*.
- xAI. 2025. Grok 4. <https://x.ai/news/grok-4>.
- Yandan Yang, Baoxiong Jia, Shujie Zhang, and Siyuan Huang. 2026. Sceneweaver: All-in-one 3d scene synthesis with an extensible and self-reflective agent. *Advances in neural information processing systems*, 38:140319–140351.
- Ruihan Zhang, Borou Yu, Jiajian Min, Yetong Xin, Zheng Wei, Juncheng Nemo Shi, Mingzhen Huang, Xianghao Kong, Nix Liu Xin, Shanshan Jiang, and 1 others. 2025a. Generative ai for film creation: A survey of recent advances. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 6267–6279.
- Yuyao Zhang, Jinghao Li, and Yu-Wing Tai. 2025b. Layercraft: Enhancing text-to-image generation with cot reasoning and layered object integration. *arXiv preprint arXiv:2504.00010*.