Produced by
AECMAGAZINE

# AI IN VISUALISATION
## TRANSFORMING ARCHITECTURE

Artificial Intelligence (AI) is having a profound impact on architectural visualisation. In this special report, learn how to turbo-charge your viz workflows with the latest AI technologies, powered by Lenovo™ ThinkStation® and ThinkPad® workstations with NVIDIA RTX™ GPUs

D5 RENDER IMAGE COURTESY OF LUIS ARNAS

Architectural visualisation has long been pivotal in bringing architectural designs to life. It enables designers and visualisation specialists to create realistic images, animations and real-time experiences of projects for better understanding, clearer communication, and compelling presentations. With the advent of artificial intelligence (AI) and AI-optimised Graphics Processing Units (GPUs) this field is undergoing a transformative evolution.

Renders that used to take minutes or hours now happen in real time. Concept design now goes beyond the sketch with generative AI producing incredible visuals to take designs in bold new directions.

All of this is made possible by AI-optimised Lenovo ThinkStation and ThinkPad P Series workstations with powerful NVIDIA RTX professional GPUs.

Rather than simply relying on 'brute force' processing, AI makes smarter use of modern compute resources.

Tensor cores in NVIDIA RTX GPUs, for example, are dedicated entirely to AI processing. They dramatically improve performance in real-time tools like Enscape, D5 Render, Chaos Vantage, NVIDIA Omniverse, and Unreal Engine, and can also slash render times in production renderers like Chaos V-Ray GPU.

NVIDIA RTX Tensor cores are also instrumental in accelerating generative AI tools like Stable Diffusion, which can produce high-quality images from textual descriptions.

In this 'AI in visualisation' special report we'll explore how software is changing, how architects can benefit from new AI-accelerated workflows and how to ensure you have the right workstation hardware in place to get the most out of these exciting new technologies.

Sponsored by Lenovo NVIDIA

# GEN AI: DESIGN INSPIRATION

Generate concept designs beyond the imagination and go from sketch to compelling visuals in record time

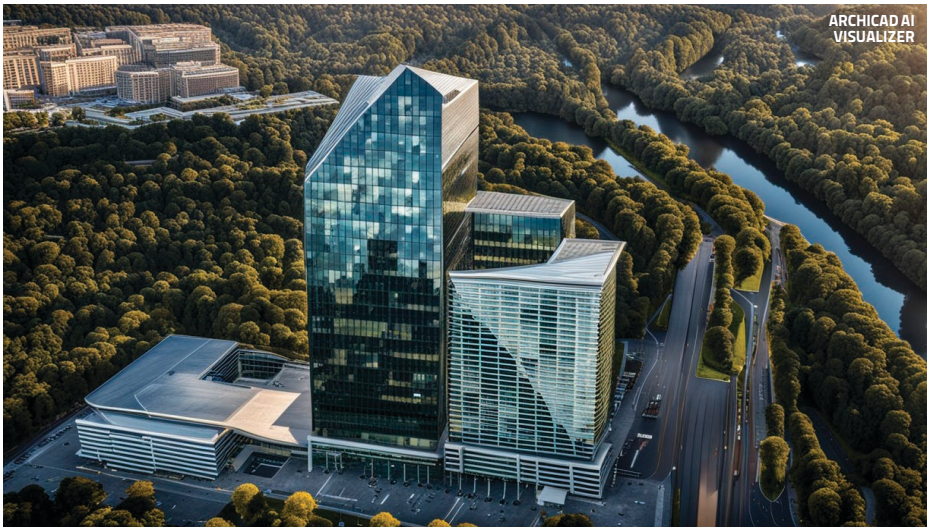


Generative AI, which involves AI models that generate text, images, and other content from the data they were trained on – burst onto the scene in 2022 and is already having a profound impact on architectural visualisation.

The most common application within architectural design is text-to-image generative AI, using software such as Stable Diffusion and Midjourney. Here an architect can quickly generate hundreds of images of architectural ideas, simply by crafting natural language descriptions, called prompts.

Text-to-image generative AI is especially useful during the early design phase, where quick exploration and iteration is critical. Traditional sketching or 3D modelling is typically constrained by the speed with which one can stroke a pen or move a mouse. With generative AI the results come back in a matter of seconds, allowing architects to explore more options than ever before.

Using AI for conceptual design also opens up new possibilities when it comes to creativity. By blending styles, design ideas generated by AI often go beyond the imagination of the architect, leading to new architectural vocabularies.

Of course, text-to-image generative AI does not come without challenges. Crafting a good prompt demands skill, and getting the desired results requires refinement and trial and error. The resulting images are often low-resolution but can be upscaled using AI.

The technology is advancing at an incredible rate. There are now new generative models and tools that provide much more control over the output.

Stable Diffusion, which is very well suited to architecture, offers several add-ons to help push output in specific directions.

The first is ControlNet, a neural network extension that allows the user to guide the AI through the composition of an uploaded image. This could be a sketch, render, photo, or even a screen grab of a concept model in SketchUp. There are several ways the AI can infer composition from the image — through depth, line art, or MLSD (Mobile Line Segment Detection), which detects straight lines, so is well suited to architecture.

The second is LoRA (Low Ranked Adaptation), a training method that is used to fine tune Stable Diffusion models by capturing the unique styles and attributes of a set of images. One could, for example, train a LoRA on a set of renders to mimic a specific architectural style.

LoRA training is very GPU intensive, and can take some time to master. The good news is there's an active community sharing LoRAs on www.civitai.com. Ismail Seleit is a London-based architect whose LoRAs include one for 'contemporary modular refined designs with realistic materials', another for architectural scale models, and even one which gives the look and feel of a hand sketch.

Open source tools like Stable Diffusion deliver excellent results, but there can be a steep learning curve. Several AEC software developers are building generative AI tools that are designed specifically for architects with a simplified set of commands that are directly accessible inside BIM software. For example, Nemetschek's AI Visualizer integrates Stable Diffusion inside Archicad, Vectorworks and Allplan. Veras is an AI renderer for SketchUp, Rhino, Revit, Forma and Vectorworks. There's also SketchUp Diffusion from Trimble. All three tools use simple 3D concept models, text prompts and sliders to guide the AI.

## GETTING STARTED WITH STABLE DIFFUSION

Stable Diffusion is one of the most popular tools for generative AI in architecture. It's not only hugely powerful, but has the added benefit of being open source. This means it's free to download and run on local hardware, so you can take full advantage of your workstation's powerful NVIDIA RTX GPU without having to pay for cloud GPUs.

To help workstation users get up and running quickly NVIDIA has published a step-by-step guide. The guide explains how to install the software and use the TensorRT extension for Stable Diffusion Web UI, using Automatic1111, the most popular Stable Diffusion distribution.

According to NVIDIA, the extension doubles the performance of Stable Diffusion by leveraging the Tensor Cores in NVIDIA RTX GPUs.

Download the guide at *www.tinyurl.com/Nvidia-SD*

# AI FOR RENDERING (DLSS)

Turbocharge your viz tools with smart AI technologies that bring new efficiencies to real-time rendering

ENSCAPE

In the realm of architectural visualisation, it's a continual challenge to strike a balance between real-time interactivity and photorealism. However, advances in Graphics Processing Unit (GPU) technology and the integration of artificial intelligence (AI) have significantly transformed this landscape.

Applications such as Enscape, Chaos Vantage, D5 Render, NVIDIA Omniverse and Unreal Engine are at the forefront of this revolution, harnessing the power of modern workstation GPUs and smart AI processing to deliver stunning visualisations with unprecedented speed and accuracy.

While generative AI may have only recently catapulted AI into the mainstream, AI has for many years played a pivotal role in advancing real-time architectural visualisation. NVIDIA Deep Learning Super Sampling (DLSS) technology has been at the forefront of this revolution.

## WHAT IS NVIDIA DLSS?

NVIDIA DLSS launched in 2018, alongside the first NVIDIA RTX GPUs, which introduced dedicated AI hardware known as Tensor Cores. DLSS has now evolved into a set of AI-accelerated technologies that software developers can integrate into their real-time visualisation tools to boost performance. This is often measured in Frames Per Second (FPS).

The exciting thing about NVIDIA DLSS is that it takes a smarter approach to graphics processing. Rather than simply throwing a bigger GPU at the problem, DLSS increases frame rates by using Tensor cores to bypass the traditional graphics pipeline with no discernible loss in visual quality. This not only means real-time experiences become smoother but the workstation can handle more complex, visually rich models.

The latest incarnation, NVIDIA DLSS 3.7 includes three distinct technologies: Super Resolution, Ray Reconstruction and Frame Generation.

DLSS Super Resolution boosts performance by using AI to output higher resolution frames from a lower resolution input. In short, one can get 4K quality output, while the GPU only renders frames at FHD resolution.

DLSS Frame Generation boosts performance by using AI to generate more frames. It only works on Ada Generation GPUs. The Tensor cores process the new frame, and the prior frame, to discover how the scene is changing, then generates entirely new frames without having to process the graphics pipeline.

DLSS Ray Reconstruction enhances image quality by using AI to generate additional pixels for intensive ray-traced scenes. It replaces hand-tuned denoisers with an NVIDIA supercomputer-trained AI network that generates higher-quality pixels in between sampled rays.

Chaos Vantage, D5 Render, Unreal Engine, and NVIDIA Omniverse were among the first real-time visualisation tools to integrate DLSS Ray Reconstruction, featuring AI-enhanced real-time preview modes with ray tracing.

According to NVIDIA, with both DLSS Frame Generation and Ray Reconstruction enabled in D5 Render, FPS in the viewport increases by 2.5x, enabling incredible resolution and visual quality in huge scenes.

Other architectural visualisation tools are using AI in different ways. Chaos V-Ray, which is renowned for its production quality output, is using NVIDIA AI denoising to decrease grainy spots and discolouration in images while minimising the loss of quality. The NVIDIA Omniverse RTX Accurate (Iray) renderer also uses AI for denoising.

D5 RENDER IMAGE COURTESY OF MARFURT

D5 RENDER IMAGE COURTESY OF LINK

# AI WORKSTATIONS FOR VIZ

Lenovo workstations with NVIDIA RTX GPUs are AI-ready and optimised for the most intensive AI tasks

Architectural visualisation has some of the most demanding workflows in AEC, especially now that AI is accelerating viz software in many different areas. In order to take full advantage of these transformative technologies — now and well into the future — it is more important than ever to have the right workstation hardware in place.

The entire range of Lenovo ThinkStation and ThinkPad P Series workstations with powerful NVIDIA RTX GPUs are 'AI-ready' and purpose built to accelerate the most challenging AI workflows. This includes a wide range of tasks for AI visualisation, as well as custom AI training, development and inferencing.

For visualisation, the Lenovo ThinkStation P8, with support for up to three NVIDIA RTX GPUs, can handle the most intensive workflows. This includes training LoRAs for Stable Diffusion or rendering V-Ray scenes at lightning speeds.

Not all architects need such powerful workstations to take advantage of AI in visualisation. The compact Lenovo ThinkStation P3 Ultra SFF, for example, provides a great entry point for generative AI and AI-accelerated real-time viz.

What's more, designers do not need to be tied to their desks. All ThinkStation workstations can be rack mounted and set up for remote access, and with Lenovo ThinkPad mobile workstations, designers can work from anywhere.

Of course, Lenovo workstations are much more than their constituent parts. Their renowned build quality is designed to withstand the rigours of professional use, while running cool and quiet, and with impressive reliability. This is critical for all viz workflows but especially for those that involve intense computation, such as training AI or rendering out 4K videos.

| Lenovo ThinkStation P3 Ultra SFF | Lenovo ThinkStation P3 Tower | Lenovo ThinkStation P8 | Lenovo ThinkPad P16 | Learn more |
|---|---|---|---|---|
| NVIDIA RTX 2000 (16 GB) or 4000 SFF (20 GB) Ada Generation GPU | NVIDIA RTX 4000 (20 GB), 4500 (24 GB) or 5000 (32 GB) Ada Generation GPU | Up to three NVIDIA RTX 6000 (48 GB) Ada Generation GPUs | NVIDIA RTX 4000 (12 GB) or 5000 (16 GB) Ada Generation laptop GPU | The Lenovo workstation portfolio |
| **AI viz workflows**<br>▪ Entry-level BIM-centric real-time viz workflows<br>▪ Image generation in Stable Diffusion | **AI viz workflows**<br>▪ Mainstream real-time viz with more complex datasets and at higher resolutions<br>▪ Image generation in Stable Diffusion | **AI viz workflows**<br>▪ High-end real-time viz and production rendering with the most demanding datasets<br>▪ Image generation and LoRA training in Stable Diffusion | **AI viz workflows**<br>▪ Entry-level to mainstream real-time viz workflows<br>▪ Image generation in Stable Diffusion | lenovo.com/aec |

## NVIDIA RTX GPUs - ACCELERATING AI VIZ WORKFLOWS

NVIDIA RTX GPUs, at the heart of Lenovo workstations, are a critical component for AI, and arch viz in general. Tensor cores within the GPUs are dedicated entirely to AI processing to massively accelerate text-to-image generative AI tools like Stable Diffusion, and dramatically boost performance in real-time viz software through NVIDIA DLSS.

NVIDIA RTX GPUs also feature RT cores for ray tracing to simulate the behaviour of light, and CUDA® cores for general tasks, including rasterisation, to turn complex viz model geometry into pixels.

For visualisation workflows, one of the most important characteristics of NVIDIA RTX GPUs is the amount of on-board memory. The NVIDIA RTX 2000 Ada Generation GPU comes with 16 GB, which is a good starting point, but for the most demanding workflows, there's the NVIDIA RTX 6000 Ada Generation with 48 GB.

GPU memory is required to load up complex viz datasets, including textures. Memory demands grow when models are viewed or rendered at higher resolutions.

Stable Diffusion is also very memory hungry, especially when training, and while there are workarounds when GPU memory is in short supply, this can have a dramatic impact on performance.

It's also important to understand that GPU memory must be shared between applications. It's not uncommon for architects to use multiple applications at the same time — render in V-Ray while working on real-time experiences in Unreal Engine, for example.

Running out of GPU memory can have a dramatic negative impact on performance, meaning users have to make compromises to the way they want to work.

Some viz workflows will benefit greatly from having more than one NVIDIA RTX GPU in a single workstation. With V-Ray GPU and NVIDIA Omniverse RTX Accurate (Iray), for example, rendering performance scales near linearly across multiple GPUs.

Of course, it's still very early days for AI in visualisation and NVIDIA RTX GPUs are already pioneering new technologies which could soon become an important part of viz workflows.

This includes Neural Rendering, which can use AI to 'learn' how light interacts with materials, instead of using ray tracing to simulate physical light behaviour, and NeRFs (Neural Radiance Fields) for adding real world context to new architectural designs.