Neural Compositing for Real-Time Graphics

6 Month Internship – La Forge, Ubisoft Bordeaux

Rendering complex visual effects, such as refractive glass, translucent materials, smoke, and other view-dependent or volumetric phenomena, often require specialized rendering techniques, including dedicated rendering passes or ray tracing. While these methods can produce high-quality results, they can also be computationally expensive, making them challenging to integrate into real-time applications under tight performance constraints.

Recent advances in neural rendering, frame generation, and generative models open new possibilities for rethinking how such effects are produced. Rather than explicitly simulating complex light transport, neural models could learn to reconstruct or approximate the final appearance of an object from its rendered representation and surrounding image context. This suggests a new rendering paradigm in which computationally expensive visual effects could be partially replaced by learned image-space transformations.

This internship will explore neural compositing, an approach in which selected objects are rendered using simplified geometric or shading proxies and subsequently enhanced by screen-space neural models that predict their final visual appearance. The objective is to investigate whether this approach can reproduce the visual complexity of effects such as refraction, translucency, and volumetric phenomena while reducing the cost of conventional rendering techniques.

The work will involve exploring recent developments in neural rendering, neural shading, and frame generation, and investigating their potential application to real-time visual effects. Different neural compositing strategies will be considered through experimental prototypes, with an emphasis on understanding how scene information and rendered inputs can be leveraged to reconstruct complex visual phenomena.

The internship is exploratory in nature, with the aim of identifying the strengths and limitations of neural compositing and determining which classes of visual effects are best suited to this approach. The findings will help assess its potential as a complement to conventional rendering techniques and inform future research on the topic.

Supervision

Antoine Houdard – antoine.houdard@ubisoft.com
Georges Nader – georges.nader@ubisoft.com

References

[1] DLSS 5: Generative Neural Rendering (https://research.nvidia.com/labs/adlr/DLSS5/)
[2] Generative AI model for Global Illumination effects (https://gpuopen.com/learn/genai-model-for-global-illumination/)
[3] Temporally stable generative illumination with a one-step diffusion model (https://gpuopen.com/learn/temporally-stable-generative-illumination/)
[4] An Introduction to Neural Shading (https://research.nvidia.com/labs/rtr/publication/duca2025neural/)