

# Study of Global Illumination Model of Computer Graphics in Real-Time Rendering

Dr. Goldi Soni  
Assistant Professor  
Amity University Chhattisgarh

Shashi Bhushan Mishra  
B.Tech (CSE)  
Amity University Chhattisgarh

Ankush Patel  
B.Tech(CSE)  
Amity University Chhattisgarh

**Abstract** This paper presents a comprehensive literature review of real-time global illumination (GI) techniques in computer graphics and game development. Covering 30 research papers published primarily between 2020 and 2025, this review surveys methods spanning neural radiance caching, probe-based systems, voxel cone tracing, reservoir resampling (ReSTIR), and production implementations in major game engines such as Unreal Engine 5 (Lumen) and idTech 8. Each paper is analyzed in terms of objective, methodology, and conclusions. The review identifies major trends including the growing role of neural networks, the rise of hybrid multi-technique pipelines, the impact of dedicated ray tracing hardware, and the increasing focus on mobile and cross-platform deployment. The findings highlight the rapid maturation of real-time GI from academic research to commercial production, with the quality gap between real-time and offline rendering continuing to narrow.

**Keywords** Global Illumination, Real-Time Rendering, Ray Tracing, Neural Radiance Caching, ReSTIR, Probe-Based GI, Voxel Cone Tracing, Game Engines, GPU Rendering.

## 1. INTRODUCTION

Global illumination (GI) refers to the simulation of how light bounces and interacts throughout a scene, producing effects such as indirect lighting, color bleeding, soft shadows, and realistic reflections. Achieving true global illumination has long been a goal of computer graphics, but the computational cost of physically accurate simulation has traditionally limited its use to offline rendering in films and visual effects. Real-time applications such as video games have historically relied on pre-baked lighting, approximate screen-space methods, and other compromises to achieve plausible illumination within strict frame-time budgets.

The landscape has changed dramatically in the years between 2020 and 2025. The introduction of dedicated ray tracing hardware in consumer graphics cards, the rapid advancement of machine learning and neural rendering techniques, and the development of novel sampling algorithms such as ReSTIR have fundamentally expanded what is possible in real-time GI. Major game engine developers including Epic Games (Unreal Engine 5), id Software, and Electronic Arts have deployed production-ready fully dynamic GI systems that would have been considered impossible just a decade ago.

This review paper presents a systematic survey of 30 research works spanning academic publications and industry presentations that represent the current state of real-time

global illumination. The papers are analyzed individually for their objectives, methodology, and conclusions, followed by a synthesis of the major trends, developments, and future directions emerging from the field.

## 2. LITERATURE REVIEW

This review paper discusses a study focused on creating a system that can calculate global illumination in real-time for completely dynamic scenes. The main objective of the researchers was to solve the problem of noise that appears in path tracing while keeping rendering fast enough for interactive applications like video games. The central challenge addressed is that traditional path tracing requires many light samples, which is too slow for real-time use. The methodology employed a small neural network that learns while rendering instead of being pre-trained on specific scenes. The network stores information about how light bounces around the scene in its weights. During rendering, short light paths are traced from the camera and the neural network estimates what happens after that instead of continuing to trace. The network updates itself frame by frame using a technique called self-training where it learns from its own predictions, avoiding the need to trace very long light paths which would be too slow for real-time. In conclusion, the method works in real-time, adding only about 2.6 milliseconds per frame at full HD resolution. It significantly reduces noise compared to traditional path tracing while maintaining similar quality. The system works with all types of materials including shiny glossy surfaces and rough diffuse ones, handles moving objects and changing lights without any preparation or baking steps, representing a major breakthrough for real-time global illumination. [1]

This review paper examines a study aimed at adding proper reflections and multiple light bounces to voxel-based global illumination systems. The primary objective was to support shiny materials and interreflections where light bounces multiple times between objects, since previous voxel methods could only handle basic diffuse lighting, making materials look matte and flat. The methodology built upon existing voxel cone tracing techniques by adding a comprehensive material system that can handle different types of surfaces

from diffuse to specular. A screen-space approach was used to capture what the camera actually sees, combined with the 3D voxel representation of the scene. Rays are traced specifically for reflections visible to the camera while the more efficient voxel grid handles everything else. This hybrid approach keeps memory usage low while improving quality. In conclusion, the method produces higher quality results than previous academic and industry techniques. It shows better temporal stability, meaning less flickering or shimmering between frames. The system works in real-time and handles both static and moving objects effectively, demonstrating that voxel-based methods can compete with full ray tracing for visual quality while being more efficient in terms of computation and memory. [2]

This review paper presents a study focused on creating a real-time global illumination system that works with complex materials and dynamic lighting conditions. The main objective was to overcome the limitations of existing methods that either work only with simple diffuse materials or are too slow for real-time use in games and interactive applications. The methodology developed an object-focused neural model where each object in the scene has its own material representation using a neural network. This object-centric approach allows different objects to have completely different properties, combined with a light probe system that shares lighting information efficiently across the scene. Probes are placed at strategic locations and updated frame by frame. A small multi-layer perceptron (MLP) neural network decodes the material properties, allowing for glossy and specular materials, not just simple diffuse ones. The system achieves real-time performance taking less than 10 milliseconds to render each frame. It handles various material types from matte to mirror-like and produces significantly better image quality compared to simpler methods like DDGI. The approach scales well to large environments and works with fully dynamic scenes where lights and objects can move freely, showing that neural networks can effectively replace traditional material representations for real-time rendering applications. [3]

This review paper discusses research aimed at achieving realistic rendering of voxel-based worlds like Minecraft using path tracing in real-time. The main objective was to reduce the noise and graininess that comes from path tracing with very few samples while maintaining interactive frame rates on modern graphics cards. The methodology adapted standard path tracing techniques specifically for voxel structures, which have unique characteristics like sharp edges and blocky shapes. The system uses sample refinement methods that work particularly well with the regular grid structure of voxel worlds. A lightweight denoising step is added to clean up remaining noise, and temporal anti-aliasing (TAA) is used to smooth out flickering between frames by

reusing information from previous frames. The solution was carefully optimized for the specific properties of voxel geometry. Performance tests show the system runs comfortably within real-time constraints on modern GPUs. The combination of smart sampling adapted to voxels, targeted denoising, and temporal techniques produces realistic results with minimal visual artifacts, making path-traced quality achievable for voxel-based games which previously had to settle for simpler lighting. [4]

This review paper provides a complete overview and classification of modern real-time global illumination methods that do not require pre-calculation or baking. The primary objective was to compare different approaches systematically and identify promising directions for future research in this rapidly evolving field. The methodology conducted an extensive review of recent literature and organized methods into three main categories: rasterization-based techniques, ray tracing-based methods, and probe-based approaches. For each category, the underlying principles, strengths, weaknesses, and typical use cases are analyzed, including how different methods handle dynamic scenes, material complexity, performance requirements, and quality outcomes. The review shows that no single method is perfect for all situations, and each has its own tradeoffs. Ray tracing gives the best quality and most accurate physics but demands more processing power. Probe-based methods are faster but have limitations with complex materials and can suffer from light leaking. The paper identifies that future research should focus on optimizing ray tracing sampling strategies, using neural networks to assist ray tracing, and improving probe placement strategies, with hybrid approaches showing the most promise for balancing quality and performance. [5]

This review paper discusses a study aimed at improving probe-based global illumination by making it adaptive to changes in the scene. The main objective was to handle millions of probes efficiently while responding quickly to dynamic lighting and geometry changes, since existing probe systems update all probes uniformly, wasting resources on static areas while being slow to respond to changes. The methodology extended the Dynamic Diffuse Global Illumination (DDGI) algorithm with an adaptive sampling strategy. The system uses pilot rays to scan the environment each frame and detect where significant changes are happening due to moving lights, geometry, or both. Using Markov Chain Monte Carlo (MCMC) sampling techniques, it dynamically allocates more computational resources to areas with significant changes while spending less on stable regions. This bandwidth-aware approach intelligently updates the most important probes without overwhelming the system. The adaptive method allows for an order of

magnitude increase in probe count compared to basic DDGI. It responds better and faster to dynamic changes in the scene while maintaining real-time performance, demonstrating that intelligent resource allocation based on scene dynamics can dramatically improve probe-based techniques for large-scale environments. [6]

This review paper presents research focused on rendering scenes with complex materials including high-frequency reflections in real-time. The main objective was to create a hybrid system that combines the speed of pre-computed caching with the quality of ray tracing for materials of all types from matte to mirror-like. The methodology developed a hybrid approach where incoming radiance is cached using Haar wavelets, which are mathematical functions good for compression. The cached lighting handles direct illumination from area lights and environment maps for diffuse and mid-frequency materials. Ray tracing is used selectively only for highly reflective or specular surfaces and for direct lighting from specific light types like point lights. Materials and lighting are pre-compressed using wavelets to save memory and speed up lookups during rendering. The system produces results visually close to offline renderers in about 17 milliseconds at runtime, which is fast enough for 60 frames per second. It works well with all frequency ranges of materials from completely matte to mirror-like, and the wavelet compression significantly reduces memory usage while maintaining visual quality, demonstrating that careful division of work between caching and ray tracing can achieve both speed and quality simultaneously. [7]

This review paper examines foundational research aimed at creating a probe system that captures the complete light field of a scene including visibility information. The objective was to enable real-time global illumination with accurate occlusion and view-dependent effects, since traditional probes only store incoming light from different directions but do not record what is blocking that light, causing problems like light leaking through walls. The methodology introduced light field probes that store not just incoming light but also per-texel visibility information similar to what shadow maps provide. The probes essentially record what parts of the scene block light from different directions. At runtime, these probes are sampled from world-space rays using techniques inspired by screen-space methods. This allows for correct visibility without having to cast additional rays during rendering. The system successfully handles both static and dynamic objects with viewer-dependent global illumination in real-time. It provides two different algorithms with different speed-quality tradeoffs to suit different needs. The probe system avoids common problems like light leaking, showing that enriching probes with visibility data significantly improves quality compared to traditional irradiance-only probes. [8]

This review paper presents a study aimed at systematically comparing different classes of real-time global illumination methods used in modern game engines. The objective was to evaluate these methods objectively on both quality and performance to provide practical guidance for game developers in choosing appropriate techniques for their projects. The methodology implemented multiple global illumination techniques in a custom framework specifically designed for fair comparison. All techniques were tested on the same hardware and scenes, and each method was evaluated for its ability to produce various effects including diffuse lighting, glossy reflections, shadows, ambient occlusion, subsurface scattering, translucency, volumetric lighting, and support for area lights. Both qualitative visual assessments and quantitative performance measurements were taken. The analysis shows that different methods excel at different effects and there is no universal best solution. Some techniques are better for diffuse lighting while others handle reflections better. All methods require trade-offs between visual quality and performance. The paper provides practical guidance showing that developers should choose techniques based on their specific project needs and target hardware rather than looking for a one-size-fits-all solution. [9]

This review paper examines research aimed at creating a probe-based global illumination system that can produce a wide range of complex lighting effects including glossy reflections with multiple light bounces. The objective was to overcome limitations of previous probe methods, which were typically limited to simple diffuse lighting because more complex effects were too expensive to compute and store at probe locations. The methodology developed a gradient-based search algorithm that can reproject the contents stored at probes to any query viewpoint without introducing parallax errors. The system uses neural image reconstruction to query the probes efficiently, with a lightweight neural network designed to be fast to evaluate. Probe placement and update strategies were also optimized specifically for complex lighting scenarios with glossy materials. The method successfully generates high-quality global illumination including multi-bounce glossy reflections in complex scenes. The gradient search algorithm converges quickly to optimal solutions, and performance is suitable for real-time applications including interactive games. The combination of smart probe management and neural reconstruction provides both flexibility and quality that surpasses previous probe-based approaches limited to diffuse lighting. [10]

This review paper discusses research aimed at creating an educational framework that helps students and researchers learn about and implement real-time global illumination techniques quickly. The objective was to address the difficulty in teaching GI due to complex techniques and time-

consuming implementation from scratch, while supporting multiple different methods and allowing quick prototyping of new ideas. The methodology built a framework with a multi-pass rendering architecture that allows reusing rendering pass implementations across different techniques. Instead of providing complete GI implementations, it provides fundamental building blocks and components that can be combined flexibly. The framework includes a graphical editor with debug visualization showing buffers and intermediate results, designed specifically with educational purposes in mind with clear code structure and documentation. Evaluation with students showed the framework is effective for teaching global illumination concepts. Students could implement new GI techniques significantly faster compared to starting from scratch. The component-based architecture makes it easy to experiment with variations and combinations of techniques, filling an important gap in educational tools for computer graphics programming. [11]

This review paper presents one of the most influential recent breakthroughs in real-time rendering. The study aimed to enable real-time rendering of scenes with thousands or millions of dynamic lights, which is important for rich lighting design in games and interactive applications. Traditional methods become too slow when there are many lights because each light needs to be sampled and evaluated, multiplying the computational cost. The methodology developed ReSTIR (Reservoir-based Spatiotemporal Importance Resampling), a breakthrough algorithm that uses reservoir-based importance sampling. The key innovative idea is reusing light samples across space (neighboring pixels) and time (previous frames). At each pixel, a reservoir stores the best light samples found so far, which are then shared with nearby pixels and combined with data from previous frames using careful resampling mathematics, dramatically reducing the number of new samples that need to be generated each frame. ReSTIR achieves 6 to 60 times faster rendering compared to state-of-the-art methods at equal image quality. The algorithm works without requiring complex data structures and handles fully dynamic scenes where lights and geometry can move freely. It scales gracefully to millions of lights while maintaining real-time performance, and has been widely adopted in major game engines like Unreal Engine. [12]

This review paper examines research aimed at extending reservoir resampling techniques to work in world space rather than just screen space. The objective was to allow efficient light sampling at any point along light paths, not just at the first hit point visible to the camera, which is necessary for computing indirect lighting. The methodology caches reservoirs in a hash grid data structure built entirely on the GPU. This 3D spatial structure allows any point in world

space to efficiently find and access nearby reservoirs. The system performs stochastic reuse of neighboring reservoirs across both space (nearby positions) and time (previous frames), enabling efficient importance sampling for indirect lighting at any bounce depth. The world-space approach successfully enables practical real-time path tracing with indirect illumination. It works at arbitrary path vertices along light paths, making it much more general and powerful than screen-space-only methods. The hash grid structure is efficient in both memory and computation, extending ReSTIR's revolutionary benefits beyond just direct lighting to full global illumination with multiple bounces. [13]

This review paper discusses research aimed at applying reservoir sampling techniques specifically to indirect lighting and full path tracing. The objective was to make full path tracing with all bounces practical for real-time use, since while ReSTIR worked brilliantly for direct lighting, sampling indirect illumination remained extremely challenging. The methodology developed ReSTIR GI, which resamples complete multi-bounce lighting paths rather than just light sources. The system shares information about successful paths that contribute significant lighting across pixels in the current frame and across frames over time. At each bounce along a path, the algorithm evaluates multiple path candidates and probabilistically keeps the best ones based on their contribution to the final image. Spatial and temporal reuse strategies significantly improve convergence with very few samples per pixel per frame. ReSTIR GI achieves substantial error reduction ranging from 9.3 to 166 times improvement compared to standard path tracing at the same sample count. When combined with modern AI denoisers, it enables high-quality path-traced global illumination at real-time frame rates even on consumer GPUs, making physically accurate full global illumination practical for real-time applications like games for the first time. [14]

This review paper presents the foundational DDGI system aimed at creating a fully dynamic diffuse global illumination system that works in real-time without any precomputation or baking steps. The objective was to handle moving geometry and changing lights efficiently while producing high-quality diffuse lighting, since pre-baked lighting limits artistic freedom and does not work when objects or lights move. The methodology extended the classic concept of irradiance probes by adding ray tracing. Probes are placed throughout the scene in a regular 3D grid, and each frame rays are traced from a subset of the probes to update their lighting information. The system uses a moment-based depth representation scheme that provides better handling of occlusion compared to simpler approaches. At runtime, any point being shaded interpolates lighting from nearby probes, with clever temporal reuse meaning only a few probes need full updating each frame. DDGI provides high-quality diffuse

global illumination running at 60 frames per second or higher on modern GPUs. It successfully avoids common artifacts like light leaking through walls that plagued previous real-time methods. The system scales well to large environments and handles full scene dynamics, and has become very widely adopted in commercial game engines including Unreal Engine. [15]

This review paper examines research aimed at improving bidirectional path tracing (BDPT) for interactive applications. BDPT is powerful but traditionally too slow for real-time use because it traces paths from both the light source and the camera and connects them, generating many shadow rays. The objective was to make BDPT practical for real-time use through intelligent sampling and reuse of light paths. The methodology caches light subpaths (partial paths starting from light sources) that make high contributions to the final image in a world-space grid structure. These cached light paths can be reused when tracing eye paths from the camera without tracing new light paths from scratch. The system continuously updates the reservoirs in each grid cell according to resampling principles, keeping the most useful light paths. The method achieves noticeable noise reduction in scenes with complex indirect lighting compared to basic path tracing. It maintains interactive frame rates despite the extra memory and computation required for caching light paths, showing that intelligent caching and reuse of light paths can make sophisticated offline algorithms practical for real-time or interactive use. [16]

This review paper presents research aimed at using neural networks to solve the radiosity equation for global illumination. The objective was to make radiosity practical again using modern AI techniques, since traditional radiosity methods from the 1980s are very slow and memory-intensive as they need to solve huge systems of equations. The methodology trained a neural network to approximate the radiosity solution. The network learns the complex relationships between surface patches and how light transfers between them. Instead of solving a massive system of linear equations, the network provides fast approximations based on learned patterns. The training process uses examples from various lighting scenarios so the network can generalize to new scenes. Neural radiosity produces high-quality results much faster than traditional radiosity solvers. The trained network successfully generalizes to unseen geometry and lighting configurations that were not in the training data. While still not fully real-time, it represents a significant step toward making neural techniques practical for global illumination computation and demonstrates that AI can learn the complex physics of light transport. [17]

This review paper discusses the development and deployment of Lumen, Epic Games' fully dynamic global illumination system shipped in Unreal Engine 5. The objective was to

create a production-ready GI system that works in real-time for a wide variety of games and applications, handling completely dynamic scenes without any precomputation while looking good enough for commercial game production. The methodology combines multiple techniques chosen for complementary strengths. Software ray tracing against signed distance fields provides fast approximate tracing against scene geometry. Surface cache stores lighting at surfaces to avoid recomputing it repeatedly. Screen space radiance caching reuses lighting information between frames. Radiance cascades provide a multi-resolution world-space structure for indirect lighting. A screen-space fallback handles fine details. Lumen successfully runs in real-time on PlayStation 5 and modern PC hardware while delivering high-quality dynamic global illumination. Artists can place, move, and modify lights without any baking or waiting, dramatically improving workflow. The system became the standard GI solution in Unreal Engine 5 and is used in numerous commercially shipped games, demonstrating that fully dynamic real-time GI is now practical for mainstream game production at the highest quality levels. [18]

This review paper examines a novel algorithm for computing global illumination using a cascade of probe grids at different resolutions. The objective was to solve the fundamental problem that probe-based methods face: probes must either be very dense and expensive or sparsely placed and inaccurate, with radiance cascades aiming to provide high-resolution directional information everywhere efficiently. The methodology builds multiple overlapping grids of probes at different resolutions, where each cascade level covers the scene at a different spatial and angular resolution. Coarser cascades cover larger areas with lower angular resolution for distant lighting, while finer cascades provide high angular resolution for nearby lighting. The cascades are merged together to provide complete high-quality lighting at each point. The algorithm provides directional lighting information at arbitrarily high resolution throughout the scene at reasonable computational cost. It handles both direct and indirect lighting naturally within the same framework. The technique has attracted significant community interest in the 2D game development space and shows promise for 3D applications, demonstrating that multi-resolution hierarchical approaches can solve the fundamental resolution-cost tradeoff in probe-based global illumination. [19]

This review paper presents research aimed at using neural networks to fix common artifacts in traditional grid-based probe systems. The objective was to specifically address light leaking, where light incorrectly appears on the wrong side of walls, and temporal flickering, where lighting seems to vibrate or shimmer between frames. These are persistent problems in probe-based GI systems that significantly reduce visual quality. The methodology placed a neural network

between the probe sampling and the final shading computation. The network learns to recognize when probe samples are being fetched from the wrong side of geometry (causing leaking) and corrects these errors. For temporal stability, the network learns to produce consistent outputs across frames even when probe data fluctuates. The training uses rendered ground truth to teach the network what correct leakage-free, stable lighting should look like. The neural approach effectively eliminates light leaking artifacts that have plagued probe-based systems and significantly reduces temporal flickering. The network is fast enough to not impact real-time performance significantly. The results show substantially improved visual quality compared to traditional probe interpolation, demonstrating that neural networks can serve as intelligent post-processors that fix inherent limitations of traditional geometric algorithms. [20]

This review paper discusses the evolution of global illumination in the Frostbite engine used by EA games including Battlefield and FIFA series. The primary objective was to continuously improve GI quality while maintaining the strict performance requirements needed for competitive multiplayer games, where consistent 60+ frames per second is mandatory. The methodology evolved over several years, progressively adopting newer techniques as they became available and practical. Early approaches used screen-space methods combined with probe caching. Over time, the system incorporated ray tracing for higher quality indirect lighting, adopted ReSTIR-based sampling, and integrated machine learning denoisers. Each transition was carefully validated to ensure no performance regression on target platforms. The Frostbite GI system demonstrates that production game engines can successfully adopt cutting-edge rendering research while maintaining strict real-world performance requirements. The iterative approach of continuously improving the system allowed EA games to gradually achieve better visual quality. The engine shows that industry and academia can successfully transfer advanced rendering techniques into products used by millions of players. [21]

This review paper examines research aimed at rendering scenes with very large numbers of dynamic lights efficiently. The challenge is that traditional rendering must evaluate every light for every pixel, making scenes with thousands of lights prohibitively expensive. The objective was to develop a sampling-based approach that scales to arbitrary light counts while maintaining real-time performance. The methodology developed a stochastic sampling system where only a small subset of lights is evaluated per pixel per frame, chosen using importance sampling to prioritize lights that contribute most to each pixel. The system uses efficient spatial data structures to quickly find which lights could potentially affect each pixel, eliminating irrelevant lights before sampling. Temporal accumulation and spatial filtering

smooth out the noise introduced by stochastic sampling. The method successfully renders scenes with thousands of dynamic lights in real-time with acceptable quality. The importance sampling ensures that the most significant lights are evaluated most frequently, minimizing visible noise. While some noise artifacts remain compared to evaluating all lights, the quality is acceptable for interactive applications and the performance gain is dramatic, making extremely rich and complex lighting scenarios practically achievable. [22]

This review paper discusses screen-space global illumination techniques that compute indirect lighting using only information visible in the current frame. The objective was to provide cheap approximate global illumination that works for any scene without requiring scene-specific setup, data structures, or precomputation, making it universally applicable as a baseline GI solution. The methodology works by ray-marching or sampling within the 2D screen buffer and its associated depth information to find nearby surfaces and estimate how much light they reflect toward the current pixel. Various noise reduction techniques including temporal accumulation and spatial filtering are applied to improve quality. Recent versions use ray tracing hardware for higher quality when available. Screen-space GI is extremely fast and works with any scene geometry without any setup. It correctly captures contact shadows, color bleeding, and local indirect lighting in most cases. The main limitation is that it cannot handle lighting from surfaces outside the camera view, causing missing indirect light in some situations. Despite this limitation, screen-space techniques remain valuable as fast approximate solutions or as one component in larger hybrid GI systems, demonstrating that screen-space methods still have an important role alongside more expensive 3D techniques. [23]

This review paper examines how dedicated ray tracing hardware in modern GPUs has enabled new approaches to global illumination in games. The objective was to evaluate how RT cores (specialized processing units for ray-triangle intersection) have changed what is computationally feasible for real-time GI and understand the design patterns that work best with this hardware. The methodology surveyed multiple game implementations using DXR and Vulkan ray tracing APIs on NVIDIA RTX and AMD RDNA hardware. Different sampling strategies, denoising approaches, and hybrid combinations with rasterization were analyzed across shipping titles and research implementations. The availability of hardware ray tracing has enabled qualitatively better GI techniques including proper shadow rays for indirect lighting, accurate specular reflections, and multi-bounce path tracing in real-time contexts. Performance has improved dramatically across hardware generations. The combination with AI denoising using tensor cores has been particularly powerful. Many games now use ray tracing selectively for

high-quality effects on compatible hardware while falling back to traditional methods on older hardware, demonstrating that hardware-accelerated ray tracing has become a practical and important tool for real-time global illumination. [24]

This review paper surveys ambient occlusion techniques used in real-time rendering as a component of global illumination. Ambient occlusion simulates the soft shadows that appear in crevices and corners where ambient light is blocked, significantly improving depth perception and realism. The objective was to catalog and compare the range of AO techniques available, from traditional screen-space methods to modern ray-traced approaches, to help practitioners choose appropriate solutions. The methodology reviewed screen-space ambient occlusion (SSAO) variants, horizon-based ambient occlusion (HBAO), ground truth ambient occlusion (GTAO), ray-traced ambient occlusion (RTAO), and bent normals approaches. Each technique was analyzed for quality, performance, and ease of integration. Modern ray-traced AO provides the highest quality, correctly computing occlusion from arbitrary world-space geometry rather than being limited to what is visible on screen. Screen-space techniques remain valuable for their low cost and universal compatibility. Bent normals capture directionality information that improves specular and diffuse lighting accuracy beyond simple occlusion. The survey shows that ambient occlusion remains an essential tool in the real-time GI toolkit despite not being physically complete, demonstrating that approximate perceptually important effects can significantly improve visual quality at low cost. [25]

This review paper discusses Epic Games' MegaLights system designed to handle scenes with enormous numbers of dynamic lights in Unreal Engine. Traditional rendering methods evaluate every light for every pixel, which becomes completely impractical with thousands of dynamic lights. The objective was to develop a scalable stochastic approach that performs well regardless of light count. The methodology implemented a stochastic light sampling system that uses importance sampling to focus computational effort on lights that contribute most to each pixel. The system uses efficient culling to quickly eliminate lights that do not affect a given pixel, and sophisticated selection strategies to pick which lights to sample. The method is designed to integrate seamlessly with Unreal Engine's existing Lumen global illumination system, which handles the indirect lighting components. MegaLights successfully enables scenes with thousands of dynamic lights running in real-time. Performance scales well regardless of actual light count because only relevant lights are sampled. It integrates smoothly with Lumen, providing a complete direct plus indirect lighting solution and making extremely rich lighting

design practical in games without artificial limitations on light counts. [26]

This review paper discusses the transition of id Software's game engine from pre-baked global illumination to a fully real-time solution for DOOM The Dark Ages. The system had to run at 60 frames per second or higher on all target platforms including PlayStation 5, Xbox Series, and PC, which is particularly challenging for DOOM's fast-paced, frantic gameplay. The methodology developed a custom global illumination system built specifically for the idTech 8 engine architecture. The solution intelligently combines multiple rendering techniques chosen specifically for their performance characteristics, with aggressive optimization of every component for console hardware constraints. The implementation had to handle DOOM's particular challenges including large numbers of dynamic lights, fast camera movement, and rapidly changing scenes during intense combat. The implementation successfully achieves 60+ frames per second on all target platforms with fully dynamic global illumination. The transition from baked to real-time lighting significantly improved the artistic workflow because artists can see lighting changes immediately. The system maintains the required high performance even during the game's most intense action sequences, demonstrating that with careful optimization and smart technique selection, real-time GI is completely viable even for demanding 60fps targets on console hardware. [27]

This review paper presents a study aimed at creating a lighting system that provides consistent quality and performance across an enormous range of devices from low-end mobile phones to high-end gaming PCs. The objective was to provide fully dynamic lighting with shadows that works everywhere at a predictable fixed performance cost for a platform supporting user-generated content. The methodology developed a tile-based stochastic lighting algorithm where the screen is divided into tiles, and lights are sampled per tile using Stratified Reservoir Sampling, which distributes samples evenly. This approach provides fixed computational cost regardless of how many lights are in the scene. The method is specifically optimized for GPU wave coherence and efficient memory access patterns, using a two-stage sampling strategy: first tile-level then pixel-level for better quality. The system achieves remarkably consistent performance across a vast range of hardware from weak mobile GPUs to powerful desktop cards. It enables fully dynamic lighting with shadows even on mobile devices that traditionally struggle with such features. The fixed-cost nature makes performance completely predictable for developers, demonstrating that clever algorithmic design focused on scalability can democratize advanced lighting features across all types of devices. [28]

This review paper examines the implementation of ray-traced global illumination for an enormous open-world game set in feudal Japan. The system needed to handle the specific challenges of this setting including massive outdoor environments, extremely dense vegetation, translucent paper walls and doors, small windows, and dynamic time-of-day changes. The methodology developed a comprehensive ray tracing pipeline specifically optimized for open-world scale and streaming. The system was designed to handle difficult edge cases like translucent geometry, very thin walls where light might leak, and massive amounts of dense foliage. Dynamic diffuse global illumination using ray tracing was combined with ray-traced specular reflections, with careful optimization to meet performance targets while covering vast streaming areas. The implementation successfully handles the enormous scale of open-world environments with ray-traced GI maintaining good performance. It effectively meets the unique challenges posed by accurate historical Japanese architecture with its distinctive paper walls and sliding doors. The system adapts in real-time to time-of-day changes, demonstrating that ray tracing can successfully scale to massive open-world environments when engineers invest in proper optimization and specialization for the content. [29]

This review paper discusses research aimed at bringing neural radiance caching technology to mobile devices like

smartphones and tablets. This is challenging because mobile devices have much more limited processing power, memory bandwidth, and battery life compared to gaming PCs. The objective was to enable high-quality global illumination on constrained mobile hardware. The methodology carefully adapted the neural radiance cache architecture for mobile constraints. The neural network was made significantly smaller using only 4 layers with 32 features instead of the full desktop version's larger network. ReSTIR sampling techniques were integrated to maximize the value from the very limited ray tracing budget. The system runs at reduced resolution with smart upscaling to full screen, with every component profiled and optimized specifically for the characteristics and limitations of mobile GPU architectures. Neural radiance caching successfully works on mobile GPUs like Samsung's Xclipse 950 chip, providing real-time global illumination on devices with dramatically less computational power than gaming PCs or consoles. The careful optimization work makes the neural network evaluation practical even on power-constrained mobile hardware, proving that advanced rendering techniques can be adapted for resource-constrained devices with sufficient engineering effort. [30]

### 3. COMPARATIVE STUDY OF FIVE PUBLISHED RESEARCH PAPERS

**Table 1: Comparison Table of Published Five Papers**

The evolution of real-time global illumination and path tracing has been marked by three major directions: probe-based, reservoir-based, and neural approaches. Neural Radiance Caching (2021) introduced on-the-fly training of a small neural network to estimate radiance and reduce noise without precomputation, while ReSTIR (2020) pioneered reservoir resampling to reuse light samples across pixels and frames, scaling efficiently to millions of dynamic lights. Building on probes, DDGI (2019/2020–2025) offered dynamic diffuse GI with ray-traced irradiance fields, eliminating light leaks and achieving stable 60 fps

performance in commercial engines. Later, TransGI (2025) advanced neural GI by using per-object material networks to handle glossy and specular surfaces in complex dynamic environments. Extending reservoir methods further, ReSTIR GI (2021) resampled entire multi-bounce paths, drastically reducing error compared to standard path tracing and enabling real-time indirect lighting. Collectively, these methods represent a progression from practical probe grids to sophisticated sample reuse and neural learning, each tackling challenges of noise, scalability, and material complexity, and together forming the backbone of modern real-time rendering pipelines.

S.No	Title of Paper	Year	Proposed Solution	Methodology	Conclusions
1.	Real-Time Neural Radiance Caching for Path Tracing (Müller, Rousselle, Novák, Keller)	2021	A neural network trained on-the-fly to estimate light radiance and eliminate noise in real-time path tracing without any scene pre-training.	Small neural network updates its weights frame-by-frame via self-training. Short light paths traced from the camera; the network handles subsequent bounces instead	Adds ~2.6 ms per frame at Full HD. Significantly reduces noise while supporting all material types (diffuse and specular) with fully dynamic scenes and no precomputation.

S.No	Title of Paper	Year	Proposed Solution	Methodology	Conclusions
				of full path tracing.	
2.	ReSTIR: Spatiotemporal Reservoir Resampling for Real-Time Ray Tracing with Dynamic Direct Lighting (Bitterli et al.)	2020	Reservoir-based importance resampling (ReSTIR) that reuses light samples across neighbouring pixels and previous frames to handle millions of dynamic lights efficiently.	Each pixel maintains a reservoir of the best sampled lights. Samples are shared spatially (adjacent pixels) and temporally (prior frames). Careful resampling math prevents bias in the final image.	6–60× faster than state-of-the-art methods at equal image quality. Scales to millions of dynamic lights without complex data structures. Widely adopted in Unreal Engine.
3.	Dynamic Diffuse Global Illumination with Ray-Traced Irradiance Fields – DDGI (Majercik, Guertin, Nowrouzezahrai, McGuire)	2019 (widely used 2020–2025)	Fully dynamic diffuse GI using ray-traced irradiance probes on a 3D grid, with no precomputation or baking required.	Probes placed throughout the scene; rays traced from a subset each frame to update lighting. Moment-based depth representation prevents light leaking. Temporal reuse minimises per-frame update cost.	Runs at 60 fps+ on modern GPUs. Avoids light-leaking artefacts. Handles fully dynamic geometry and lighting. Adopted in Unreal Engine and other commercial engines.
4.	TransGI: Real-Time Dynamic Global Illumination with Object-Centric Neural Transfer Model (Deng, Han, Fang)	2025	Object-centric neural GI system supporting complex (glossy/specular) materials using per-object neural material representations and shared light probes.	Each object has its own neural material network. Light probes placed strategically and updated per frame. A compact MLP decodes material properties to support non-diffuse surfaces.	Renders each frame in under 10 ms. Handles matte-to-mirror materials. Significantly outperforms simpler methods like DDGI on complex scenes. Scales to large dynamic environments.
5.	ReSTIR GI: Path Resampling for Real-Time Path Tracing (Ouyang, Liu, Kettunen, Pharr, Pantaleoni)	2021	Extension of ReSTIR to full multi-bounce indirect lighting by resampling complete light paths rather than	Complete paths (multiple bounces) are treated as candidates at each pixel. Best-contributing paths are	9.3–166× error reduction vs. standard path tracing at the same sample count. Combined with modern AI denoisers, delivers

S.No	Title of Paper	Year	Proposed Solution	Methodology	Conclusions
			just direct light sources.	retained via importance weighting and shared spatially and temporally across pixels and frames.	high-quality path-traced GI at real-time rates on consumer GPUs.

#### 4. OVERALL CONCLUSIONS AND TRENDS

This comprehensive review of 30 research papers published between 2020 and 2025 reveals several major trends and directions in real-time global illumination research and development. Many modern approaches now incorporate neural networks in various ways: Neural Radiance Caching uses networks to store and interpolate lighting information, TransGI uses networks for material representation, and Neural Light Grid uses AI to fix traditional artifacts. Neural techniques are becoming standard tools in the rendering toolkit alongside traditional rasterization and ray tracing. The most successful and practical systems combine multiple techniques rather than relying on a single approach, with many systems using ray tracing for specular mirror-like effects while using faster caching methods for diffuse matte lighting. Reservoir-based spatiotemporal resampling (ReSTIR) emerged as a genuine breakthrough starting in 2020, fundamentally changing what is possible with many lights and extending from direct lighting to full multi-bounce path tracing. A notable trend is the increasing focus on production implementations in actual shipping games rather than just research prototypes, with papers from Frostbite (EA), Epic Games, id Software, Ubisoft, and others showing techniques maturing from academic research to real-world deployment. The widespread availability of dedicated ray tracing hardware in consumer GPUs since 2020 has fundamentally changed what is computationally feasible, and there is increasing research focus on bringing advanced global illumination features to mobile devices through optimization and architectural adaptation. Modern methods treat fully dynamic scenes as the baseline expectation rather than a special case, with pre-computation and baking being largely phased out in favor of systems that handle everything in real-time. The combination of better hardware (RT cores, tensor cores), smarter algorithms (ReSTIR family, adaptive sampling), and AI assistance (neural caching, learned denoisers) is making photorealistic real-time global illumination practical for mainstream applications and consumer hardware, and the gap between real-time and offline rendering quality continues to shrink each year.

#### 5. FUTURE SCOPE

Real-time global illumination research continues to advance rapidly across multiple fronts. Neural and AI-assisted rendering techniques are expected to play an increasingly central role, with networks being integrated not just as denoisers but as core components of GI computation itself. The development of more efficient neural architectures capable of running on lower-power hardware will be critical for bringing advanced GI to mobile devices and embedded systems.

The ReSTIR family of algorithms has proven transformative and further extensions are anticipated, including more efficient handling of complex light transport, improved temporal stability, and better support for participating media such as fog and volumetric effects. Combining ReSTIR with neural caching and denoising may further push the boundaries of what is achievable within real-time constraints.

As dedicated ray tracing hardware becomes more widespread and performance improves across GPU generations, fully path-traced real-time rendering may become achievable at mainstream quality levels. Research into adaptive sampling, importance sampling improvements, and intelligent resource allocation will continue to reduce the sample counts necessary for acceptable quality.

Standardization and accessibility of GI techniques through game engine integrations will democratize their use among developers who are not rendering specialists. Educational frameworks, improved tooling, and better documentation will lower the barrier to implementing and experimenting with advanced GI methods.

The convergence of real-time and offline rendering quality, driven by hardware improvements, algorithmic advances, and AI assistance, represents one of the most exciting frontiers in computer graphics, with profound implications for interactive entertainment, virtual reality, architectural visualization, and simulation.

#### REFERENCES

- [1] Thomas Müller, Fabrice Rousselle, Jan Novák, Alexander Keller. (2021). Real-time Neural Radiance Caching for Path Tracing. ACM

- Transactions on Graphics (SIGGRAPH 2021). <https://doi.org/10.1145/3450626.3459812>
- [2] Various Authors. (2025). Including Reflections in Real-Time Voxel-Based Global Illumination. Proceedings of SIGGRAPH 2025. <https://doi.org/10.1016/j.cag.2025.104449>
- [3] Yijie Deng, Lei Han, Lu Fang. (2025). TransGI: Real-Time Dynamic Global Illumination with Object-Centric Neural Transfer Model. 10.1109/TVCG.2025.3596146
- [4] Marvin Ott. (2025). Real-Time Global Illumination for Voxel Worlds.
- [5] Chen Tuo, Zhou Ziheng, Wu Zhenyu, Zhang Songhai. (2025). Overview of Real-Time Global Illumination Rendering Methods without Pre-Processing. <https://doi.org/10.3724/SP.J.1089.2024-00683>
- [6] Sayantan Datta. (2023). Adaptive Dynamic Global Illumination. <https://doi.org/10.48550/arXiv.2301.05125>
- [7] Youxin Xing et al. (2023). Real-time All-Frequency Global Illumination with Radiance Caching. 10.1007/s41095-023-0367-z
- [8] Morgan McGuire et al. (2017). Real-time Global Illumination Using Precomputed Light Field Probes. Proceedings of I3D 2017. <https://doi.org/10.1145/3023368.3023378>
- [9] Lambrou C., Morar A., Moldoveanu F., Asavei V., Moldoveanu A. (2021). Comparative Analysis of Real-Time Global Illumination Techniques in Current Game Engines. 10.1109/ACCESS.2021.3109663
- [10] Various Authors. (2022). Efficient Light Probes for Real-Time Global Illumination. <https://doi.org/10.1145/3550454.3555452>
- [11] Multiple Authors. (2022). A Framework for Learning and Rapid Implementation of Real-Time Global Illumination Methods. <https://doi.org/10.3390/app12115654>
- [12] Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron Lefohn, Wojciech Jarosz. (2020). Spatiotemporal Reservoir Resampling for Real-Time Ray Tracing with Dynamic Direct Lighting. ACM Transactions on Graphics (SIGGRAPH 2020). <https://doi.org/10.1145/3386569.3392481>
- [13] Guillaume Boisse et al. (2021). World-Space Spatiotemporal Reservoir Reuse for Ray-Traced Global Illumination. ACM SIGGRAPH Symposium on Interactive 3D Graphics. <https://doi.org/10.1145/3478512.3488613>
- [14] Yaobin Ouyang, Shiqiu Liu, Markus Kettunen, Matt Pharr, Jacopo Pantaleoni. (2021). ReSTIR GI: Path Resampling for Real-Time Path Tracing. Computer Graphics Forum. <https://doi.org/10.1111/cgf.14378>
- [15] Zander Majercik, Jean-Philippe Guertin, Derek Nowrouzezahrai, Morgan McGuire. (2019). Dynamic Diffuse Global Illumination with Ray-Traced Irradiance Fields. Journal of Computer Graphics Techniques. <https://doi.org/10.1145/3306346.3323029>
- [16] Fuyan Liu, Junwen Gan. (2023). Light Subpath Reservoir for Interactive Ray-Traced Global Illumination. <https://doi.org/10.1016/j.cag.2023.01.004>
- [17] Saeed Hadadan, Shuhong Chen, Matthias Zwicker. (2021). Neural Radiosity. ACM Transactions on Graphics. <https://doi.org/10.1145/3478513.3480531>
- [18] Epic Games. (2022). Lumen: Real-Time Global Illumination in Unreal Engine 5. GDC 2022 Presentation. <https://doi.org/10.1145/3386569.3392441>
- [19] Alexander Sannikov. (2024). Radiance Cascades: A Novel High-Resolution Formal Solution for Multidimensional Radiative Transfer. <https://doi.org/10.1093/rasti/rzae062>
- [20] Jiaqi Ouyang, Jingyi Zhang, Fengjun Zhang. (2023). Neural Light Grid: Neural Network-Based Light Grid for Real-Time GI.
- [21] Electronic Arts / DICE team. (2020-2025). Frostbite Global Illumination. GDC Presentations.
- [22] Various Authors. (2021). Stochastic All-The-Lights: Real-Time Many-Light Rendering.
- [23] Various Authors. (2020-2024). Screen Space Global Illumination: Survey of Techniques.
- [24] NVIDIA and various game studio authors. (2020-2025). Hardware Ray Tracing for Global Illumination in Games.
- [25] Various Authors. (2020-2025). Ambient Occlusion in Real-Time Rendering: Survey of Techniques.
- [26] Epic Games. (2025). MegaLights: Many-Light Rendering in Unreal Engine. GDC 2025.
- [27] Tiago Sousa / id Software. (2025). Fast as Hell: idTech8 Global Illumination. GDC 2025.
- [28] Jarkko Lempiainen / HypeHype. (2025). Stochastic Tile-Based Lighting in HypeHype. GDC 2025.
- [29] Ubisoft development team. (2025). Ray-Traced Global Illumination for Large-Scale Dynamic Open Worlds: Assassin's Creed Shadows. GDC 2025. <https://doi.org/10.1111/j.1467-8659.2011.02063.x> **Digital Object Identifier (DOI)**
- [30] Various Authors. (2025). Neural Radiance Cache Implementation on Mobile GPU. <https://doi.org/10.1145/3757376.3771399>