Monte Carlo (MC) integration is an essential tool in many fields of science. In image synthesis, it has enabled photorealistic rendering results via physically based light transport simulation. However, an inherent problem of MC integration is variance causing noise in rendered images. This thesis presents three methods, each taking a different approach to variance reduction in rendering.

The first approach focuses on improving the sampling technique. An adaptive solution is proposed for unbiased direct illumination sampling, employing Bayesian regression and a novel statistical model of direct illumination to achieve robustness. This method was integrated into a production renderer, demonstrating both its theoretical soundness and practical utility.

The second approach explores the combination of multiple sampling techniques via multiple importance sampling (MIS). Optimal weighting functions are derived, proving to minimize the variance of MIS estimators. The new weights outperform all common heuristics and provide novel design considerations for selecting appropriate sampling techniques in integration problems.

Finally, the third approach involves pre-computation to handle challenging scenarios effectively. Pre-computed reference images of a clear sky are used to create a high-quality fitted model, allowing any renderer to achieve realistic sky appearance without any atmospheric simulation overhead. An extension covering the full spectral range of terrestrial solar irradiance is also introduced, enabling usage of such pre-computed models for purposes other than renderings intended to mimic the perception of human observers, such as thermal analysis, and photovoltaic plant yield simulations.

The three presented methods significantly improve rendering efficiency and quality, and contribute valuable insights to the field of MC integration in image synthesis.