8 Further Graphics (aco41)

This question is on rendering a scene via ray tracing and the rendering equation.

(*a*) You are tasked with rendering a video of a scene via ray tracing. Only the camera is moving, the rest of the scene is static, the light sources emit light of a single wavelength, and surfaces do not emit light. You are allowed to shoot one ray per pixel and do pre-computation for a rotation and position of the camera. For each of the following cases, explain if this is possible, and if so, describe what you would pre-compute and store in single-channel textures, and how you would finally render the scene.

(*i*) Local illumination and a single BRDF (bidirectional reflectance distribution function) for all surfaces. [3 marks]

(*ii*) Global illumination and diffuse reflection for all surfaces. [3 marks]

(*b*) We are rendering a scene via path tracing and approximating integrals with importance-sampling.

(*i*) Assume you can importance-sample with the exact incoming radiance. Write the resulting approximation of outgoing radiance at any point and direction with N samples. What is the length (number of segments) of the light path to compute this approximation? Why is this length finite? [5 marks]

(*ii*) There is only a single point light source in a scene, only light paths of length three are allowed, and no shadow rays are allowed. What would be your importance-sampling strategy, and why at the first and second intersection points for a given path? [4 marks]

(*c*) For a given scene, assume the rendering equation can be reduced to $L_o(\mathbf{x}, \vec{\omega}_o) = c \int_{H^2} V(\mathbf{x}, \vec{\omega}_i) \cos\theta_i d\vec{\omega}_i$ at surface points $\mathbf{x}$, where V is the (binary) visibility function.

(*i*) What are the properties of the scene that make this reduction possible? [3 marks]

(*ii*) Assuming a single object and ignoring inter-reflections for that object, what is the gradient of $L_o(\mathbf{x}, \vec{\omega}_o)$ with respect to $\mathbf{x}$? [2 marks]