Carnegie Mellon University

The Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213

September 14, 2023

Dear committee,

I am writing as a referee for the thesis submitted by **Tomáš Iser** titled **Affordable optical measurement methods for predictive rendering**. This is an area within computer graphics, and more broadly physics and optics, that I have expertise in, having published multiple works on similar topics in the past. After reviewing the thesis, I have found that it introduces several technical contributions of potential importance to not just computer graphics, but also other scientific and application areas. I will elaborate on these points in the rest of this referee report. Overall, my assessment is that the thesis successfully demonstrates the candidate's ability to carry out creative scientific work, and should be sufficient for the completion of his doctoral degree.

The thesis falls within the broader area of *inverse rendering*. The high-level statement of the inverse rendering problem is as follows: We assume we are given a set of image measurements of some scene for which we are missing some information – say, the shape, pose, or material properties of objects in it. Then, we want to infer from these images the missing information, so that we can use a digital twin of the scene to simulate, or *render*, synthetic images that match the image measurements. Hence, this procedure acts as a sort of *inverse* to the rendering process (instead of going from a scene to images as in rendering, we are going from images to a scene), hence the name. Once the missing information is inferred, it can be used to not only match input image measurements, but also synthesize other scene images, or for further analysis tasks (3D processing, material analysis, and so on). This is a very broad problem statement, which encompasses several subareas within not only computer graphics, but more generally computer vision, optics, and applied physics.

Within this broader context, this thesis focuses on inverse rendering for optical properties of two important classes of materials: translucent and fluorescent materials. Translucent materials are those that *scatter* light at their interior (a processed known as *subsurface scattering*), rather than (only) reflect it at their surface. These materials are ubiquitous, including important classes of materials such as all biological tissues, clouds, smoke, fog, and chemicals such as soap, pigments, and inks. Fluorescent materials are those that absorb light of one wavelength, and reemit it at a different wavelength. These materials are likewise commonplace, especially among synthetic materials, such as pigments, and inks. Because of the complex light-matter interaction processes that characterize these materials, solving inverse rendering problems for their optical properties is a challenging task that has been researched in several sciences, including computer graphics, remote sensing, applied physics, medical imaging, and material sciences.

For each of these two inverse rendering tasks, the thesis makes twofold contributions. First, it proposes an optical acquisition system, that is design to be: 1. discriminative, in the sense that the image measurements it provides can discriminate between different values for the optical properties of the underlying material; 2. affordable, in the sense that the optical apparatus can be assembled at low cost in terms of both hardware and engineering effort. Second, it proposes a computational post-processing system that can invert the acquired image measurements and accurately recover the unknown optical material properties, in a way that correctly accounts for the complex light-matter interactions that characterize translucent and fluorescent materials.

Even though I describe the two contributions as distinct, it is important to emphasize here that successfully putting together the end-to-end inverse rendering pipeline requires *jointly* designing the optical acquisition and computational post-processing systems. For example, any acquisition system can only be calibrated to some level of accuracy. Therefore, the design of the acquisition system should be such that miscalibration errors are of a magnitude that is within the tolerance levels the subsequence postprocessing system can tolerate. Likewise, different types of acquisition systems result in post-processing systems that have different computational complexities – some types of imaging measurements are more challenging to render than others – in a way that is typically inversely related to the information content in the acquired measurements. Therefore, one needs to carefully consider the trade-offs between information content, calibration accuracy, and computational complexity, before arriving at a successful design.

On the acquisition system side, the thesis proposes two systems that, compared to prior art, prove remarkably effective: Not only do they provide verifiably – as I will describe next – accurate measurements, but also they are significantly cheaper to build *and* calibrate compared to prior systems. On this aspect I can speak from personal experience, as designing an acquisition pipeline for translucent materials was indeed the focus of my thesis. I can tell you first-hand how expensive and difficult it was to design, assemble, and calibrate the system my thesis proposed! By contrast, this thesis proposes a much simpler system, and explains in detail how one would go about calibrating it. Likewise, the same is true for the acquisition of fluorescent materials.

On the post-processing system side, the thesis proposes two systems that leverage recent advances in inverse rendering in the form of *differentiable rendering*: This is a type of optical simulation that, instead of images, synthesizes *derivatives* of images with respect to scene parameters – such as optical scattering and fluorescence parameters. The reason why this is significant is that one can combine these derivatives with a gradient-based optimizer to efficiently solve an inverse rendering problem. These differentiable rendering pipelines were originally proposed specifically for inverse rendering of scattering parameters – in my own thesis – and have recently received a lot of attention in computer graphics, resulting in new algorithms of much increased generality and performance.

Putting everything together, this thesis shows results focusing on a specific type of materials, namely inks and pigments used for conventional and 3D printing. The results show very accurate acquisition, particularly with respect to the spectral-dependency of the optical properties of these materials. Having accurate models for printing inks and pigments is critical for 3D printing and fabrication applications, which mix calibrated inks to reproduce different types of appearance. In fact, there is a whole industry dedicated to devices for characterizing optical properties of 3D printing inks, none of which can achieve, to the best of my knowledge, the accuracy of the methods proposed in this thesis.

More broadly, the ability to acquire optical properties of translucent and fluorescent materials accurately and efficiently is of potential application to many other scientific and application areas. Examples include particle sizing (characterizing industrial polydispersions), medical imaging (characterizing tissue samples and phantom materials), remote sensing (characterizing soil samples), oceanic studies (characterizing water samples), and so on. Therefore, I have no doubt that the results of the proposed thesis have potential for strong impact, not only within computer graphics, but broadly in sciences and applications alike.

As a concluding note, the thesis is an example of truly *interdisciplinary* research, bridging together theory and techniques from optics, rendering, and material modeling. Successfully completing the inverse rendering systems shown in the thesis is a rare accomplishment, given that it requires very tedious and technically involved work in all of these areas. Together with the important of the problem and the potential of impact, it is my overall assessment that this thesis contributes significant research outcomes. Taking these points into account, I have no doubt that the thesis proves the author's ability to produce creative scientific work.

Please do not hesitate to contact me with any questions.

Sincerely,

Ioannis Gkioulekas
Associate Professor
The Robotics Institute
School of Computer Science
Carnegie Mellon University
Contact: 1-617-840-9778, igkioule@cs.cmu.edu