Banff International Research Station Workshop on Computational Light Transport

Todd Zickler, Kiriakos Kutulakos, Gordon Wetzstein (Organizers)

1 Introduction

It is easy to forget, from casual observation, that the speed of light is finite. Our eyes are simply not fast enough to resolve the quadrillionth-of-a-second separations that exist between the refraction, reflection, and scattering events that a photon endures along its path from source to photoreceptor.

Imagine what happens when these femtosecond photon events can be resolved, and when detected photons can effectively be "tagged" with information about their starting location, the length of the path they travelled, and the events they experienced along the way. This could transform atmospheric science and aerial navigation by allowing complete three-dimensional reconstructions of cloud volumes. It could disrupt biology and medicine by enabling non-invasive probing of deep biological tissues. It could revolutionize robotics through the creation of sensors that can unambiguously analyze everyday scenes according their material properties and three-dimensional shapes, even when these scenes are beyond the line of sight or obscured by fog or smoke.

Somewhat shockingly, the dawn of this future has already arrived. It has been brought about by the sudden confluence of many factors that have all matured at the same time: new sensors; new programmable light sources; more powerful computing engines; and new, data-driven algorithmic tools. At the heart of all this is a young new field of interdisciplinary study, called Computational Light Transport, that uses emitters, sensors, and computation to control and analyze the flow of light.

This workshop brought together experts from the various backgrounds that are involved in computational light transport—including optics, computer vision, and computer graphics—to help solidify a scientific community in this young area, and to identify key developments, challenges, research directions, and opportunities. This workshop was especially important because many of these experts come from different academic communities and do not otherwise have opportunities to meet as a complete group.

1.1 Why this topic at BIRS?

Computational Light Transport is the synergistic convergence of three areas: computer vision, computer graphics, and applied optics. Over the past decade, this convergence has produced imaging systems with unprecedented capabilities, including the ability to:

- capture videos of light-in-flight at billions or trillions of frames per second;
- tomographically reconstruct turbid 3D volumes, even when there is substantial off-axis scattering;
- capture images around corners by processing photons that travel beyond the line of sight; and
- by using "compressed measurements", capture images using far fewer samples than is dictated by the Nyquist sampling rate.

At the core of these developments, and at the core of Computational Light Transport, is the seamless fusion of optics and computation. It is the creation of new cyber-physical systems that combine digital computations with physical ones, the former in the conventional manner, and the latter by generating specialized lighting patterns and manipulating light on its way to a photosensor. The key insight is that by designing the cyber and physical aspects of computation together, one can create visual sensors that have transformative capabilities.

It is entirely possible that these recent successes of Computational Light Transport are just the tip of the iceberg, and that the field is on the cusp of explosive development that will bring it through to maturity. However, the field currently lacks the mathematical foundations that this maturation would require. In the new realm of Computational Light Transport many conventional models, bounds, and limits no longer apply. This includes bounds on imaging resolution, as illustrated by the compressed imaging example above, as well as limits on the reconstructability of geometric and radiometric scene properties. In order for the field to properly mature, new bounds and limits must be established. And for this to happen, we must first identify some mathematical, statistical, and computational models that can properly unify the cyber and physical aspects of Computational Light Transport.

Although a substantial number of international research groups are working in areas related to Computational Light Transport, many of them publish in different scientific journals and present at different conferences. Thus, there is currently no single venue that allows these diverse researchers to meet as a complete group in order to network, form a community, and identify the grand challenges.

The fundamental goal of this workshop was to help solidify a scientific community in the area Computational Light Transport, and to identify unifying mathematical models and representations that will support future theoretical analyses and provide design tools for the engineers of the future.

1.2 Why now?

The year 2019 marks a crucial moment in the development of Computational Light Transport. There are several reasons.

First, the proliferation of materials and devices for controlling all aspects of detection and illumination—including temporal characteristics, phase, intensity, wavelength, and polarization—has dramatically changed what physical information we can sense and how we can optically probe for otherwise invisible visual information via coded illumination. For example, novel sensors such as time-resolved single-photon detectors together with ultra-short pulsed lasers, have enabled unprecedented imaging modalities such as seeing through scattering media and around corners via inverse light transport analysis.

Second, an expansion of rapid prototyping capabilities for opto-mechanical components and devices has made it increasingly more convenient to explore novel optical coding strategies for engineering Computational Light Transport. A variety of spatial light modulators make it possible to optically code any dimension of high-dimensional light transport. Also, 3D printing, injection moulding, laser cutting, soft lithography, and other fabrication technologies have made it possible to quickly prototype complex mechanical parts at low cost.

Third, deep learning and other data-driven inference techniques and algorithms are explosively redefining many scientific fields, including imaging and computer vision. Fueled by the rapidly increasing availability of large-scale datasets suitable for learning statistics of geometry, materials, and plenoptic functions in natural scenes, convolutional neural networks and other machine learning architectures have signaled the beginning of a new era in imaging and optical scene reconstruction.

Finally, high-performance computing resources, such as cloud computing systems and graphics processing units (GPUs), have become ubiquitous during the last few years. High-end GPUs today process hundreds of gigaFLOPS (floating point operations per second), and distributed cloud-computing services place almost unlimited processing power within range of everyone. Similarly, the software tools for programming GPUs and cloud systems have become significantly more user friendly, such that even undergraduate students can easily solve large-scale problems in the cloud.

2 Topics

The workshop consisted of ten sessions. As a collection, these sessions were guided by five basic aims: (i) summarizing existing mathematical models and bounds; (ii) summarizing the present and near-term capabilities of sensing and optics technologies; (iii) identifying the mathematical, statistical, and signal processing challenges that are most readily within reach; (iv) identifying long-term open research problems and grand challenges; and (v) building consensus around unifying mathematical and statistical representations of shape, optical materials, and light transport. Each of the ten sessions is summarized below.

2.1 Ultra-fast Detectors

Session leaders: Edoardo Charbon (Ecole polytechnique federale de Lausanne), Joyce Poon (University of Toronto), Roman Genov (University of Toronto)

Session summary: Single-photon detectors, in particular gated sensors and single-photon avalanche diodes (SPADs) have grown from a curiosity to a multi-billion dollar market, with proximity sensors and LiDARs. The state-of-the-art in SPADs and SPAD image sensors reached 1/4 megapixels in size and in performance, with 30ps timestamping and a few picosecond gating. These imagers have enabled several applications relevant to this community, including around-the-corner vision, time-in-flight vision, and specific vision modes, such as fluorescence lifetime, and multi-spectral vision.

An alternative to SPAD technology is the use of conventional CMOS a image sensors with *ad hoc* pixel-level gating techniques that enable non-conventional aperture schemes. Current progress and future directions for CMOS image sensors with pixels that are individually programmable for per-pixel coded exposure, with multiple buckets that store sorted photo-generated charge is an interesting development in this field achieving significant traction with the computational imaging community.

Finally, the emergence of silicon integrated photonics, a technology enabling the mass-manufacturing of nanophotonic components has created new opportunities in generation, modulation, and detection of light, as well as ongoing research extending the technology into the near infrared and visible spectrum. We can now explore how integrated photonics will be used to create unique light sources and sensors for computational imaging.

2.2 Ultra-fast Imaging

Session leaders: Andreas Velten (University of Wisconsin), Edoardo Charbon (Ecole polytechnique federale de Lausanne), Mohit Gupta (University of Wisconsin), Sanjeev Koppal (University of Florida)

Session summary: A variety of emerging computational imaging methods rely on fast illumination and transient light detection on timescales of nanoseconds or picoseconds. This session provided an introduction to fast light detectors with a focus on Single Photon Avalanche Diodes (SPAD) in the context of several emerging application areas. It introduced the state of the art of SPADs along with their limitations and future potential. It then reviewed three areas of current and emerging applications: Non-Line-of-Sight imaging, High-Dynamic Range Imaging, and Structured Illumination Imaging.

Non-Line-of-Sight (NLOS) imaging uses transient illumination information captured from a relay surface together with an inverse light transport model to create reconstructions of scenes from the point of view of the relay surface. Time of flight NLOS techniques crucially rely on the ability to illuminate and image with high speeds and the capabilities of NLOS systems are directly connected to the capabilities and trade-offs provided by the capture hardware. The session reviewed the current use of pulsed light sources and gated SPAD detectors and provided some perspectives for future detector development and its implications for the future of NLOS imaging.

Recent research has also shown the potential of SPADs for High Dynamic Range Imaging. While a conventional camera pixel aims to measure the total number of incoming photons over a fixed exposure time, a SPAD based sensor can estimate the incoming flux by measuring timing statistics of the incoming photon stream. This has the potential to provide much higher dynamic range than conventional methods.

2.3 Lensless Imaging

Session leaders: Wolfgang Heidrich (King Abdullah University of Science and technology), Ori Katz (The Hebrew University of Jerusalem), Ramesh Raskar (Massachusetts Institute of Technology), Laura Waller (UC Berkeley)

Session summary: This session featured 5 speakers on a broad spectrum of imaging systems that don't use lenses, as well as lively discussions on when and where lensless imaging becomes useful, and how to quantify performance for these non-traditional cameras that use computation to form an image from raw non-image data. We discussed mask-based lensless camera made from a scattering element in front of a sensor, as well as multiple methods for seeing around corners and through scattering material, by using the random scattering to generate spatio-angular diversity. We also discussed how to estimate movement and position of people walking around behind a corner, by analyzing the shadows on the ground. We considered applications in military devices, medical imaging, and photography, with methods ranging from simple signal processing to advanced machine learning.

2.4 Coherent Computational Imaging

Session leaders: Jason Fleischer (Princeton University), Ori Katz (The Hebrew University of Jerusalem), Laura Waller (Uc Berkeley), Changhuei Yang (California Institute of Technology)

Session summary: Prof. Fleischer began the session by introducing the basics of correlations and their behavior in both the near and far fields (described by the van Cittert-Zernike theorem). He then demonstrated shaping of spatial correlations using a spatial light modulator. This coherence engineering gives non-traditional statistics, useful for both basic research and as a degree of freedom for imaging with partially incoherent light. Finally, he discussed the propagation of statistical light in a nonlinear medium, in which the coherence evolves as a function of its intensity and modal distribution. Nonlinear coupling between modes allows interaction between signal and noise, which under certain parameter ranges can result in a stochastic resonance that amplifies the signal at the expense of the noise.

Prof. Waller discussed the modeling and application of speckle fields. She compared ray vs. wave optics and showed when the two descriptions could coincide, e.g. at small defocus for partially developed caustics. She then showed how speckle correlation imaging, via the convolution model of the memory effect, breaks for thicker diffusers and higher scattering angles. Usefully, she showed how one can model variations from a perfect shift using a few Principle Components with a low-rank model.

Prof. Katz made use of temporal coherence to demonstrate time-gated speckle correlations and passive time of flight. In the former, he showed result on low-coherence interferometry that localized measurement to a thin slice along the optical axis; transverse imaging then followed from direct rays or speckle correlations. In the latter, he used the low coherence to ensure phase matching between optical path lengths (from the source to the scattering plane to the receiver), in a manner similar to passive radar. He demonstrated localization in all three dimensions, for multiple point sources. In turn, these sources could be used as reference points for other methods, such as speckle correlation imaging that is agnostic to depth.

2.5 Computer-in-the-loop Imaging

Session leaders: Ren Ng (UC Berkeley), Matthew O'Toole (Carnegie Mellon University), Srinivasa Narasimhan (Carnegie Mellon University), Achuta Kadambi (UC Los Angeles)

Session summary: This session began with deep-dive presentations about several niche applications involving computation-in-the-loop imaging. These served as case studies of the broad topic, and then the session opened up to a general discussion about technical opportunities for computation in computational imaging.

The first niche application was presented by Ren Ng on imaging and tracking the human retina in real-time at cellular scale, then stimulating individual photoreceptor cells in order to create various novel percepts of color. The computation-in-the-loop was of a fairly garden-variety conceptually (albeit with state-of-the-art performance requirements), but the specific applications in neuroscience and color perception are unprecedented and were of high general interest. The second niche application was presented by Matthew O'Toole on closed-loop feedback techniques for computational illumination and sensing. These techniques consist of adaptive, scene-dependent sensing or illumination strategies used to analyze scenes more efficiently. Applications discussed include optical tone mapping, light transport acquisition, and hyperspectral imaging.

The third niche application was presented by Achuta Kadambi on Artificial Physics: bringing imaging physics to the deep learning revolution. The computation-in-the-loop arises in the incorporation of physical priors into the design of the neural networks. A landscape of different strategies to combine priors with networks was presented.

The general discussion was opened up and led by Srinivasa Narasimhan.

2.6 Connecting microscopic and macroscopic scattering

Session leaders: Roarke Horstmeyer (Duke University), Matthew O'Toole (Carnegie Mellon University)

Session summary: Optical scattering continues to make it difficult to capture clear images in all kinds of settings. Whether it is trying to image through haze or fog, capture light off a semi-rough surface or peer into tissue, scattering constantly comes up as something that we need a better understanding of. This session explored the bridge between the types of scattering models that exist for macroscopic imaging (e.g., through fog) and those that exist for microscopic imaging (e.g., through tissue). For both scales, we discussed the transmission matrix approach to modeling the light transport. We also discussed examples of using this matrix to form the forward problem and the inverse problem. Finally, we discussed example properties of this matrix and how they relate to effects like limited scattering and the optical memory effect.

2.7 Challenges & Opportunities for Underwater Imaging

Session leaders: Tali Treibitz (University of Haifa), Mohit Gupta (University of Wisconsin), Andreas Velten (University of Wisconsin), Achuta Kadambi (Uc Los Angeles)

Session summary: The ocean covers 70% of the earth surface, and influences almost every aspect in our life, such as climate, fuel, security, and food. The ocean is a complex, vast foreign environment that is hard to explore and therefore much about it is still unknown. Interestingly, only 5% of the ocean floor has been seen so far and there are still many open marine science questions. All over the world depleting resources on land are encouraging increased human activity in the ocean, for example: gas drilling, desalination plants, port constructions, aquaculture, fish farming, producing bio-fuel, and more. As human access to most of the ocean is very limited, most of the operations in it rely on remote sensors, many of which do not exist yet or do not provide sufficient information. Thus, the future calls for substantial related research. The uncertainty stems from the fact that most existing technology cannot be applied in the ocean, as the ocean poses numerous challenges in ocean engineering such as handling optics through a medium, movement, limited resources, communications, power management, and autonomous decisions, while operating in a large-scale environment.

The panel in the workshop started with a review of the importance of optical imaging in the ocean, and continued with discussing the current state-of-the-art and the challenges in passive and active systems. Then, we discussed the feasibility and challenges of optics-based passive navigation in the sea. Such systems are highly needed because of lack of GPS signal underwater. However, the research in this field is very much in its infancy.

2.8 Optical Imaging Deep Inside the Body

Session leaders: Srinivasa Narasimhan (Carnegie Mellon University), Ramesh Raskar (Massachusetts Institute of Technology, Ioannis Gkioulekas (Carnegie Mellon University)

Session Summary: This session was aimed at creating a moonshot project for the research community—developing an optical camera that can image deep below several layers of tissue at cellular scale, without any incision. If we succeed, this technology could span a range of practical uses, ranging from novel wearables to new non-invasive point-of-care that could revolutionize healthcare. The beauty of our stated goal is that with each additional mm of depth we achieve, we can actually enable many new clinical and healthcare applications, ensuring continuous success.

But this problem is too hard to be solved by any single field so the session's aim was to bring together researchers from a wide variety of areas including computational imaging, sensor design, biomedical optics, machine learning and inverse graphics. Virtually all the attendees of this workshop could have an important role in this project. In order to streamline the discussions, the session leaders presented an overview of the problem, the current state of the art in optical imaging through tissue and the limitations. More importantly, they provided a framework to address this moonshot with three strongly connected thrusts: (a) photon path cameras and sources that should have unprecedented control over the probabilistic paths that photons take beneath the skin, (b) estimating clinically relevant bio-optical properties of tissue and vasculature by making breakthroughs in inverse graphics computations, and (c) building actual devices using this architecture that address critical practicalities such as allowable form-factor, energy consumption and operating time by distributing computations across optical, analog and digital domains. Srinivasa Narasimhan presented this framework and raised several questions for discussion. Ramesh Rasker discussed the promise of doing multi-modal sensing (using sound, light and heat). Finally, Ioannis Gkioulekas discussed the promise of developing differentiable rendering methods for addressing inverse light propagation through tissue. The session ended with Kiriakos Kutulakos discussing the idea of optimizing computations for this problem.

2.9 Bridging the Gap Between Physical Optics Propagation and Physically-based Rendering

Session leaders: Marc Christensen, Ioannis Gkioulekas (Carnegie Mellon University), Jason Fleischer (Princeton University), Laura Waller (Uc Berkeley), Andreas Velten (University of Wisconsin), Anat Levin (Technion)

Session summary: Current tools for optical design provide limited functionality when it comes to simulating field transport (irradiance + phase) in complex macroscopic scenes comprised of assortment of surface types (optically rough). Physically based graphics renderers provide a partial solution to the problem by simulating steady state irradiance transport at macroscopic scales. In recent times, the capability of these renderers have been extended to accommodate spatially and temporally patterned illumination sources. Despite advances, the ability to propagate phase in complex macroscopic scenes remains an elusive problem. The principal challenge lies in identifying computationally efficient approaches to simulating light scatter (irradiance + phase) off optically rough surfaces to produce physically plausible speckle that exhibits well-documented properties such as the angular/spectral correlation of scattered fields, and the Van-Cittert Zernike theorem.

This session was organized as a panel that engaged the workshop participants in identifying approaches and instruments that will help develop new physically based renderers that explicitly accommodate the wave nature of light and model the wide range of coherence effects observed in practice. Such renderers have wide ranging applications including holography, optical metrology and the emerging area of indirect imaging.

2.10 The Future of Imaging Systems

Session leaders: William Freeman (Massachusetts Institute of Technology), Sara Abrahamsson (University of California Santa Cruz)

Session summary: In this session, Prof. Sara Abrahamsson presented her work on structured illumination aberration-corrected multi-focus microscopy (MFM-SIM, Abrahamsson et al. BOE 2017) using computationally designed diffractive Fourier optical elements. The major challenge in using diffractive optics in wide-field imaging is challenges of correcting for chromatic dispersion, with two solutions presented by Dr. Abrahamssom: single sensor systems with a combination of one diffractive and one refractive optical element, and a new method from her lab (unpublished) where a multi-camera system can be built in a manner that allows for chromatic dispersion correction with a single diffractive element. This solution is in the SaraLab currently being employed in fast, high-resolution fluorescence microscopy — with immediate applications in Cell Biology and Neurobiology imaging — but could also be of general interest in other fields of optical imaging such as photography and astronomy. The following discussion revolved around how to improve the computational design of custom optical elements but also of how to computationally design new illumination strategies that could optimize the image system performance in challenging scenarios where the microscopic light transport includes scattering and wavelength-dependent effects. This discussion now may lead to a collaboration in exploring these between the SaraLab and computational labs present at the conference.

In the second part of this session, Prof. Bill Freeman presented some of his group's recent work on passive non-line-of-sight imaging, or "looking around corners". In this work, the low-intensity signals of ambient light reflected off of objects hidden from direct line of sight that impinge on a directly visible floor within the line of sight of a camera. Bill and his group developed algorithms that estimate the motion of the hidden objects from these subtle variations of this indirectly reflected light. Bill introduced the concept of using penumbras as a more general light transport concept that could be used to recover the 3D shape of directly visible and also hidden objects. He proposed that this technique may even be useful to calculate the 3D shape of the surface of the moon from observed shadows. This also led to discussion about the field's engagement in science education and outreach to the public.

3 Participants

Abrahamsson, Sara	UC Santa Cruz
Bouman, Katie	Caltech
Charbon, Edoardo	EPFL
Christensen, Marc	Southern Methodist University
D'eon, Eugene	unaffiliated
Davis, James	UC Santa Cruz
Fleicher, Jason	Princeton
Freeman, Bill	MIT
Genov, Roman	University of Toronto
Gkioulekas, Ioannis	Carnegie Mellon University
Goyal, Vivek	Boston University
Gupta, Mohit	University of Wisconsin at Madison
Heide, Felix	Princeton University
Heidrich, Wolfgang	KAUST
Horstmeyer, Roarke	Duke University
Kadambi, Achuta	MIT
Katz, Ori	Hebrew University
Koppal, Sanjeev	University of Florida
Kutulakos, Kyros	University of Toronto
Levin, Anat	Technion University
Narasimhan, Srinivasa	Carnegie Mellon University
Ng, Ren	UC Berkeley
O'Toole, Matthew	Stanford University
Poon, Joyce	University of Toronto
Rangarajan, Prasanna	Southern Methodist University
Raskar, Ramesh	MIT
Schechner, Yoav	Technion University
Swedish, Tristan	MIT
Treibitz, Tali	University of Haifa
Velten, Andreas	University of Wisconsin at Madison
Waller, Laura	UC Berkeley
Wetzstein, Gordon	Stanford University
Yang, Changhuei	California Institute of Technology
Zickler, Todd	Harvard University